RUGGEDCOM®
INDUSTRIAL STRENGTH NETWORKS™

ROX®
Rugged Operating System™
on Linux

v2.2 Web Interface User Guide

For RuggedRouter® RX1000

January 30, 2012

ROX™: Web Interface User Guide

Copyright © 2011 RuggedCom Inc.

ALL RIGHTS RESERVED

Dissemination or reproduction of this document, or evaluation and communication of its contents, is not authorized except where expressly permitted. Violations are liable for damages. All rights reserved, particularly for the purposes of patent application or trademark registration.

This document contains proprietary information, which is protected by copyright. All rights are reserved. No part of this document may be photocopied, reproduced or translated to another language without the prior written consent of RuggedCom Inc.

Disclaimer Of Liability

We have checked the contents of this manual against the hardware and software described. However, deviations from the description cannot be completely ruled out.

RuggedCom shall not be liable for any errors or omissions contained herein or for consequential damages in connection with the furnishing, performance, or use of this material.

The information given in this document is reviewed regularly and any necessary corrections will be included in subsequent editions. We appreciate any suggested improvements. We reserve the right to make technical improvements without notice.

Registered Trademarks

RuggedServer™, RuggedWireless™, RuggedCom Discovery Protocol™ (RCDP™), RuggedExplorer™, Enhanced Rapid Spanning Tree Protocol™ (eRSTP™), ROX™, Rugged Operating System On Linux™, RuggedBackbone™ are trademarks of RuggedCom Inc. Rugged Operating System® (ROS®) and RuggedSwitch® are registered trademarks of RuggedCom Inc. Other designations in this manual might be trademarks whose use by third parties for their own purposes would infringe the rights of the owner.

Warranty

Five (5) years from date of purchase, return to factory. For warranty details, visit www.ruggedcom.com or contact your customer service representative.

Contacting RuggedCom

Corporate Headquarters	US Headquarters	Europe Headquarters
RuggedCom Inc. 300 Applewood Crescent Concord, Ontario Canada, L4K 5C7 Tel: +1 905 856 5288 Fax: +1 905 856 1995 Toll-free: 1 888 264 0006	RuggedCom 1930 Harrison St., Suite 209 Hollywood, Florida USA, 33020 Tel: +1 954 922 7938 ext. 103 Fax: +1 954 922 7984 Toll-free: 1 888 264 0006	RuggedCom Unit 41, Aztec Centre, Aztec West, Almondsbury, Bristol United Kingdom BS32 4TD Tel: +44 1454 203 404 Fax: +44 1454 203 403
Email: RuggedSales@RuggedCom.com		

Technical Support
Toll Free (North America): 1 866 922 7975 International: +1 905 856 5288 Email: Support@RuggedCom.com

Web: www.RuggedCom.com

Table of Contents

Preface 19
Supported Platforms 19
Who Should Use This User Guide 19
How This Guide Is Organized 19
Applicable Operating System Software Revision 19
I. Administration 20
1. The ROX™ Web Interface 21
1.1. Getting Started 21
1.1.1. Requirements 21
1.1.2. Connecting To The Web Interface 21
1.1.3. The Web Browser Connection 21
1.2. The Structure Of The Web Interface 22
1.2.1. Top-level Menu Categories 24
1.3. Making Configuration Changes 25
1.3.1. Configuring Tables Using Key Settings Forms 26
1.3.2. Viewing More Information in Tables 28
2. System Administration 30
2.1. Administration menu 30
2.2. System Commands 30
2.3. Administrative Access Control 34
2.4. User Accounts 38
2.5. Software Upgrade 40
2.6. ROXflash Cross-Partition Imaging Tool - Software Downgrade 43
2.6.1. Uses 43
2.6.2. ROXflash Configuration 43
2.7. Scheduling Jobs 45
2.8. The Featurekey 48
2.8.1. Overview 48
2.8.2. Upgrading Feature Levels in the field 48
2.8.3. When a File-based featurekey does not Match the Hardware 48
2.8.4. Viewing RuggedCom Serial Numbers 49
2.8.5. Uploading a Featurekey 50
2.8.6. Backing Up a Featurekey Using the Web User Interface 51
2.9. Installing and Backing Up Files 52
2.9.1. Installing Files 52
2.9.2. Backing Up Files 53
2.10. Deleting Log Files 54
2.11. Saving Full Configurations 54
2.12. Loading Full Configurations 55
3. Time Synchronization 56
3.1. NTP Fundamentals 56
3.1.1. The NTP Sanity Limit 56
3.2. Configuring Time Synchronization 57
3.2.1. Configuring the System Time and Date 57
3.2.2. Configuring the System Time Zone 57
3.2.3. Configuring the Local Time Settings 58
3.2.4. Configuring NTP Servers 58
3.2.5. Adding Server Keys 60
3.2.6. Configuring NTP Server Restrictions 60
3.2.7. Configuring an NTP Server using Multicast or Broadcast 62
3.2.8. Configuring an NTP Client using Multicast 63

3.2.9. Configuring an NTP Client using Broadcast 63
3.2.10. Checking NTP Status 64
4. Basic Network Configuration 65
4.1. IP Interfaces 65
4.1.1. Configuring an IP Address 65
4.1.2. Simple Network Setup with the Default IPv4 Addresses 66
4.1.3. Configuring an IPv6 Address 67
4.1.4. Simple Network Setup with IPv6 Addresses 67
4.1.5. Routable Interfaces 68
5. IP Network Interfaces 70
5.1. IPv6 Fundamentals 70
5.1.1. Addressing 70
5.1.2. Security 70
5.1.3. IPv6 Address Scopes 70
5.1.4. IPv6 Multicast Addresses 70
5.2. IPv6 Neighbor Discovery 71
5.3. Non-switched Interface Menu 74
5.3.1. Configuring IP Address Source and ProxyARP for Non-switched Interfaces 76
6. Alarms 79
6.1. Introduction 79
6.1.1. Alarm Subsystems 79
6.1.2. Fail-Relay Behavior 79
6.1.3. Alarm LED Behavior 79
6.1.4. Clearing and Acknowledging Alarms 79
6.2. Alarm Configuration 80
6.2.1. Administrative Alarm Configuration 83
6.2.2. Chassis Alarm Configuration 84
7. Domain Name Search 85
7.1. Domain Name Lookup 85
8. Logging 86
8.1. Configuring Local Syslog 86
8.2. Configuring the Remote Syslog Server 86
8.3. Deleting Logs 89
9. SNMP 90
9.1. SNMP Traps 90
9.2. SNMP Access Configuration 92
9.2.1. Add an SNMP User ID 92
9.2.2. Create an SNMP Community 93
9.2.3. Map the Community to a Security Group 94
9.3. SNMP menu 94
9.4. SNMP Discovery 98
9.5. SNMP Community 98
9.6. SNMP Target Addresses 99
9.7. SNMP Users 101
9.8. SNMP Security to Group Maps 103
9.9. SNMP Access 103
10. Authentication 106
10.1. RADIUS 106
10.1.1. RADIUS overview 106
10.1.2. RADIUS Usage 106
10.1.3. RADIUS on ROX™ 107
10.1.4. RADIUS, ROX™, and Services 107
10.1.5. RADIUS Authentication Configuration 107

11. NETCONF 110
12. Chassis Management 114
12.1. Power Controller 115
12.2. Slot Hardware 116
12.3. Slot Identification 117
12.4. CPU 118
12.5. Slot Status 119
12.6. Slot Sensors 120
12.7. Module Configuration 121
13. PPP Users 124
13.1. Overview 124
13.2. PPP Configuration 124
13.3. PPP Interfaces and Link Failover 127
14. DHCP Server 129
14.1. DHCP Fundamentals 129
14.1.1. DHCP Network Organizations 129
14.1.2. Option 82 Support with Disable NAK 129
14.2. Configuring DHCP Server 130
14.2.1. Enabling the DHCP Service 130
14.2.2. DHCP Interfaces 130
14.2.3. DHCP Subnets and Pools 131
14.2.4. DHCP Shared Networks 132
14.2.5. DHCP Hosts 132
14.2.6. DHCP Host-groups 133
14.2.7. Viewing Active DHCP Leases 133
14.2.8. DHCP Options 134
14.2.9. Custom DHCP Options 139
14.2.10. Hardware Configuration 139
II. Network Interfaces and Ethernet Bridging 141
15. Ethernet Statistics 142
15.1. Viewing Non-switched Ethernet Statistics 142
16. IP Statistics 146
17. Virtual Switch Bridging 149
17.1. Overview 149
17.1.1. Helpful Hints 149
17.2. Sample Use Case 150
17.3. Virtual Switch Configuration and Status 151
18. Modem 157
18.1. Modem Configuration and Status 157
18.2. PPP and the Embedded Modem 160
18.2.1. PPP and Modem Fundamentals 160
18.3. PPP and the Cellular Modem 170
18.3.1. PPP and Cellular Modem Fundamentals 170
18.3.2. PPP Cellular Modem Information and Configuration 171
19. Serial Protocols 185
19.1. Introduction 185
19.1.1. Serial IP Port Features 185
19.1.2. Serial Protocols Applications 186
19.1.3. Serial Protocols Concepts And Issues 187
19.1.4. TcpModBus Server Application 188
19.1.5. TcpModbus Concepts And Issues 188
19.1.6. DNP (Distributed Network Protocol) 191
19.2. Serial Protocol Configuration 192
19.2.1. Assigning Protocols 192

19.2.2. Setting Rawsockets 195
19.2.3. Setting TcpModbus 196
19.2.4. Setting DNP 198
19.3. Serial Protocol Statistics 199
19.3.1. Transport Connections 201
19.4. Restarting the Serial Server 203
19.5. Resetting Ports 203
20. WAN 204
20.1. T1/E1 Fundamentals 204
20.1.1. Frame Relay 204
20.2. T3/E3 Fundamentals 204
20.3. WAN Configuration 205
20.3.1. T1 Parameters 206
20.3.2. E1 Parameters 206
20.3.3. T3/E3 Parameters 207
20.3.4. E3 Parameters 208
20.3.5. Configuring Protocols 208
20.3.6. Loopback Test 216
20.4. Statistics 217
20.4.1. Physical Layer-related Statistics 217
20.4.2. Protocol-related Statistics 223
20.4.3. Clearing Statistics 229
20.5. Location Of Interfaces And Labeling 229
20.6. LED Designations 230
20.7. DDS 230
20.7.1. DDS Configuration 230
20.7.2. Viewing and Clearing DDS Statistics 234
III. Routing and Security 236
21. Tunnelling 237
21.1. IPsec 237
21.1.1. VPN Fundamentals 237
21.1.2. IPsec Configuration 240
21.2. L2TP Tunnelling Configuration 250
21.3. Layer 2 Tunnelling 252
21.3.1. IEC61850 GOOSE Fundamentals 252
21.3.2. Generic Layer 2 Tunnel Fundamentals 253
21.3.3. Layer 2 Tunnelling Configuration 254
21.4. Generic Routing Encapsulation (GRE) 262
21.4.1. Generic Routing Encapsulation Configuration 262
22. Dynamic Routing 265
22.1. Introduction 265
22.1.1. RIP, OSPF, and BGP 265
22.1.2. RIP Fundamentals 265
22.1.3. OSPF Fundamentals 265
22.1.4. Key OSPF And RIP Parameters 266
22.1.5. OSPF And VRRP Example Network 268
22.1.6. BGP Fundamentals 270
22.2. Dynamic Routing Configuration 270
22.3. RIP 270
22.3.1. RIP Configuration 271
22.4. OSPF 276
22.4.1. OSPF Configuration 277
22.5. BGP 281
22.5.1. BGP configuration 281

23. Static Routing 288
24. Routing Status 290
24.1. IPv4 290
24.2. IPv6 291
24.3. Memory Statistics 291
24.4. RIP 293
24.5. OSPF 293
24.6. BGP 298
25. Multicast Routing 301
26. Firewall 305
26.1. Firewall Fundamentals 305
26.1.1. Stateless vs Stateful Firewalls 305
26.1.2. Linux® netfilter, iptables, and the Firewall 305
26.1.3. Network Address Translation 305
26.1.4. Port Forwarding 306
26.2. Firewall Quick Setup 306
26.3. Firewall Terminology And Concepts 307
26.3.1. Zones 307
26.3.2. Interfaces 307
26.3.3. Hosts 308
26.3.4. Policy 308
26.3.5. Masquerading and SNAT 309
26.3.6. Rules 310
26.4. Configuring The Firewall And VPN 311
26.4.1. Policy-based Virtual Private Networking 311
26.4.2. Virtual Private Networking to a DMZ 312
26.5. Firewall Configuration 312
26.5.1. Adding a Firewall 313
26.5.2. Working with Firewall Configurations 314
26.5.3. Zone Configuration 315
26.5.4. Interface Configuration 316
26.5.5. Host Configuration 317
26.5.6. Policies 318
26.5.7. Network Address Translation 319
26.5.8. IP Masquerading 320
26.5.9. Rules 321
27. Traffic Control 325
27.1. Traffic Control Modes 325
27.1.1. Traffic Control Basic (basic-configuration) Configuration Mode 325
27.1.2. Traffic Control Advanced (advanced-configuration) Configuration Mode 325
27.2. Traffic Control Configuration 327
27.2.1. Traffic Control Modes 327
28. VRRP 344
28.1. VRRP Fundamentals 344
28.1.1. The Problem With Static Routing 344
28.1.2. The VRRP Solution 344
28.1.3. VRRP Terminology 344
28.2. VRRP Configuration 346
28.2.1. VRRP Status 349
29. Link Failover 351
29.1. Path Failure Discovery 351
29.2. Using Routing Protocols and the Default Route 351
29.3. Configuring Link Failover 352

29.3.1. Configuring the Link Failover Settings 352
29.3.2. Setting a Link Failover Backup Interface 354
29.3.3. Setting a Link Failover Ping Target 354
29.3.4. Link Backup On Demand 355
29.3.5. Viewing Link Failover Status 355
29.3.6. Viewing the Link Failover Log 356
29.3.7. Testing Link Failover 356
IV. Appendices 358
A. Upgrading Software 359
A.1. Preparing The Software Upgrade 359
A.2. Launching The Upgrade 360
A.3. Monitoring The Software Upgrade 361
B. RADIUS Server Configuration 365
B.1. PPP / CHAP and Windows IAS 365
C. Setting Up An Upgrade Server 366
C.1. Upgrade Server Requirements 366
C.2. Initial Upgrade Server Setup 366
C.3. Upgrading The Repository 366
C.4. Setting Up The Routers 367
C.5. Using Microsoft Internet Information Services (IIS) Manager 6.0 or Higher as a ROX Upgrade Repository 367
D. GNU General Public License 368
D.1. Preamble 368
D.2. TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION 369
D.2.1. Section 0 369
D.2.2. Section 1 369
D.2.3. Section 2 369
D.2.4. Section 3 370
D.2.5. Section 4 370
D.2.6. Section 5 370
D.2.7. Section 6 370
D.2.8. Section 7 371
D.2.9. Section 8 371
D.2.10. Section 9 371
D.2.11. Section 10 371
D.2.12. NO WARRANTY Section 11 372
D.2.13. Section 12 372
D.3. How to Apply These Terms to Your New Programs 372

List of Figures

1.1. The ROX™ Login Form 22
1.2. The ROX™ Web Interface 22
1.3. Top-level Menu 24
1.4. Example of Edit Private Mode 25
1.5. Adding Key Information 27
1.6. Key Information in a Table 27
1.7. Example of Key Settings 1 28
1.8. Example of Key Settings 2 28
1.9. First Table of Information 29
1.10. Second Table of Information 29
2.1. Administration menu 30
2.2. Clear All Alarms Menu Action form 30
2.3. Acknowledge All Alarms Menu Action form 30
2.4. Shutdown the Device Menu Action form 31
2.5. Reboot the Device Menu Action form 31
2.6. Set New Time and Date form 31
2.7. Set Clock on Target Device form 31
2.8. Restore-factory-defaults Trigger Action form 32
2.9. Administration form 32
2.10. Hostname form 32
2.11. Timezone form 33
2.12. Setting the Timezone Form - in Edit Private Mode 33
2.13. Current System Time form 33
2.14. CLI Sessions form 34
2.15. Idle-timeout field 35
2.16. Session Limits form 35
2.17. STFP Sessions form 36
2.18. WWW Interface Sessions 37
2.19. Idle-timeout field 38
2.20. Users menu 38
2.21. Users table 38
2.22. Users form 39
2.23. Users Screen in Edit Private View 39
2.24. Software-Upgrade menu 40
2.25. Upgrade Settings 40
2.26. Upgrade Monitoring 41
2.27. Launch Upgrade 42
2.28. Decline Upgrade 42
2.29. Rollback and Reboot 42
2.30. ROX-Imaging menu 43
2.31. ROXflash Monitoring form 44
2.32. ROXFlash menu 44
2.33. ROXFlash forms 45
2.34. Scheduler menu 45
2.35. Scheduled-jobs table 46
2.36. Scheduled Jobs Form 46
2.37. CLI in the ROX™ Web Interface 49
2.38. Install Files forms 50
2.39. Backup Files forms 52
2.40. Administration menu 52
2.41. Install Files forms 53

2.42. Backup Files forms 53
2.43. Delete-logs menu 54
2.44. Delete Log Files form 54
2.45. Save-full-configuration menu 54
2.46. Save Full Configuration forms 55
2.47. Load-full-configuration menu 55
2.48. Load Full Configuration forms 55
3.1. Set new Time and Date form 57
3.2. Timezone form 58
3.3. Local Time Settings form 58
3.4. Network Time Protocol (NTP) Servers 58
3.5. Network Time Protocol (NTP) Servers form 59
3.6. Server Keys form 60
3.7. Server Restrictions Key settings form 61
3.8. Server Restrictions form 61
3.9. NTP Broadcast/Multicast Servers form 62
3.10. NTP Multicast Clients form 63
3.11. Network Time Protocol (NTP) form 63
3.12. NTP Service Status form 64
4.1. IP menu 65
4.2. Configuring an IP Address 66
4.3. Basic Network Setup Using the Default IPv4 Addresses 66
4.4. Simple IPv6 Network Setup 68
4.5. Routable Interfaces table 68
4.6. Routable Interfaces form 68
4.7. Addresses table 69
4.8. Addresses form 69
5.1. Neighbor Discovery form 72
5.2. Neighbor Discovery IPv6 Prefix 73
5.3. Neighbor Discovery IPv6 Prefix forms 73
5.4. Non-switched Interface menu 74
5.5. Routable Ethernet Ports table 74
5.6. Routable Ethernet Ports form 75
5.7. Configuring Dynamic Address Source and ProxyARP 77
6.1. Alarms menu 80
6.2. Active Alarms table 80
6.3. Active Alarms Key Settings form 80
6.4. Active Alarms form 81
6.5. Clear Alarm Menu Action form 82
6.6. Acknowledge Alarm Menu Action form 82
6.7. Clear All Alarms Menu Action form 82
6.8. Acknowledge All Alarms Menu Action form 82
6.9. Admin Alarm Configuration table 83
6.10. Admin Alarm Configuration form 83
6.11. Chassis Alarm Configuration table 84
6.12. Chassis Alarm Configuration form 84
7.1. DNS menu 85
7.2. Domain Name Searches form 85
7.3. Domain Name Servers 85
8.1. Logging menu 86
8.2. Remote Server table 86
8.3. Remote Server form 87
8.4. Remote Server Selector table 87
8.5. Selector menu 87

8.6. Remote Server Selector form 88
9.1. Adding an SNMP User ID 92
9.2. Creating an SNMP Community 93
9.3. Mapping the Community to a Security Group 94
9.4. SNMP menu 94
9.5. SNMP Sessions form 95
9.6. SNMP USM Statistics form 97
9.7. SNMP-Discover action 98
9.8. SNMP Engine ID Discover forms 98
9.9. SNMPv1/v2c Community Configuration table 98
9.10. SNMPv1/v2c Community Configuration form 99
9.11. SNMP Target Configuration table 99
9.12. SNMPv3 Target Configuration form 100
9.13. SNMP User Configuration table 101
9.14. User Configuration Key Settings form 102
9.15. SNMP User Configuration form 102
9.16. SNMP Security Model to Group Mapping table 103
9.17. Key Settings form 103
9.18. SNMP Security Model to Group Mapping form 103
9.19. SNMP Group Access Configuration table 104
9.20. Key Settings form 104
9.21. SNMP Group Access Configuration form 104
10.1. Authentication menu 106
10.2. Primary RADIUS Server form 108
10.3. Secondary RADIUS Server form 108
11.1. NETCONF menu 110
11.2. NETCONF Sessions form 110
11.3. Idle-timeout field 111
11.4. NETCONF State/Statistics form 112
12.1. Chassis menu 114
12.2. Chassis Status form 114
12.3. Power Controller form 115
12.4. Power Status table 115
12.5. Power Status form 115
12.6. Slot Hardware table 116
12.7. Slot Hardware form 116
12.8. Slot Identification table 117
12.9. Slot Identification form 117
12.10. Slot CPU/RAM Utilization table 118
12.11. Slot CPU/RAM Utilization form 118
12.12. Slot Status table 119
12.13. Slot Status form 119
12.14. Slot Sensors table 120
12.15. Slot Sensors form 120
12.16. Modules table 121
12.17. Modules form 121
12.18. Fixed Modules form 122
12.19. Fixed Modules table 122
12.20. Module Database table 123
12.21. Module Database form 123
12.22. Configurable Modules table 123
12.23. Configurable Modules form 123
13.1. PPP menu 124
13.2. Dial-in PPP Users table 124

13.3. Dial-in Users form 125
13.4. Dial-out PPP Users table 125
13.5. PPP Configuration form 125
13.6. PPP Primary Radius Server form 127
13.7. PPP Secondary Radius Server form 127
14.1. DHCP Server menu 130
14.2. DHCP Server form 130
14.3. Listen Interfaces table 131
14.4. Subnet Configuration table 131
14.5. Subnet Configuration form 131
14.6. IP Pool Configuration table 132
14.7. Shared Network Configuration table 132
14.8. Host Configuration table 132
14.9. Host Group Configuration table 133
14.10. /services/dhcpserver/show-active-leases form 134
14.11. Lease Configuration form 135
14.12. Client Configuration form for Subnets and Shared Networks 135
14.13. Client Configuration form for Hosts 136
14.14. Client Configuration form for Host-groups 136
14.15. Client Configuration form for DHCP Clients 137
14.16. NIS Configuration form 138
14.17. Netbios Configuration form 138
14.18. Setting a DHCP Custom Option 139
14.19. Hardware Configuration form 139
15.1. Statistics Menu 142
15.2. Routable-Only Ethernet Port Status Form 142
15.3. Receive Statistics Form 144
15.4. Transmit Statistics Form 144
16.1. Interfaces IP Menu 146
16.2. Routable Interface Statistics Table 146
16.3. Routable Interface Statistics Form 146
16.4. Receive Statistics Form 147
16.5. Transmit Statistics Form 147
17.1. Virtual switch with multiple interfaces 150
17.2. Adding a Virtual Switch 151
17.3. Interface Virtualswitch menu 151
17.4. Virtualswitch table 151
17.5. Virtualswitch form 152
17.6. Interface table 152
17.7. VLAN table 152
17.8. VLAN form 153
17.9. Interfaces Virtualswitch menu 153
17.10. Virtualswitch table 153
17.11. Virtualswitch form 153
17.12. Receive form 154
17.13. Transmit form 154
17.14. VLAN table 155
17.15. VLAN Receive form 155
17.16. VLAN Transmit form 156
18.1. Modem menu 157
18.2. Routable Modem Interfaces table 157
18.3. Routable Modem Interfaces form 157
18.4. Interfaces Modem menu 158
18.5. Modem Interfaces table 158

18.6. Interfaces Modem submenu 159
18.7. Modem Interfaces form 159
18.8. PPP Interfaces Statistics form 159
18.9. Reset Trigger Action 160
18.10. At Command and Trigger Action forms 160
18.11. Interface Modem submenu 160
18.12. ip/mdm-slot-1 161
18.13. Configuring a PPP Profile for Dial-out 162
18.14. PPP Configuration form 163
18.15. Configuring a PPP Profile for Dial-in 164
18.16. Dial-in PPP Users table 164
18.17. Dial-in Users form 164
18.18. V90 menu 165
18.19. V90 PPP Client Configuration 165
18.20. V90 PPP Server Configuration form 166
18.21. V90 Modem Configuration form 166
18.22. V90 Modem form 167
18.23. V90 PPP Server Configuration form 167
18.24. External Modem Configuration form 168
18.25. External Modem form 169
18.26. External PPP Server Configuration form 169
18.27. Interfaces Cellmodem menu 171
18.28. HSPA Cellular Modem Information form 171
18.29. Edge Cellular Modem Information form 173
18.30. Global Cellular GSM menu 174
18.31. GSM Cellular Network Configuration form 174
18.32. PPP Configuration form 175
18.33. CDMA EVDO Cellular Modem Information form 176
18.34. CDMA Over The Air Activation form 177
18.35. CDMA Over The Air Activation Trigger Action form 178
18.36. CDMA Manual Activation form 178
18.37. CDMA Manual Activation Trigger Action form 178
18.38. CDMA Reset Modem Trigger Action form 179
18.39. Global Cellular CDMA menu 179
18.40. Cellular Network Configuration table 179
18.41. Cellular Network Configuration form 179
18.42. PPP Configuration form 180
18.43. Interface Cellmodem menu 181
18.44. Routable Cellular Modem Interfaces table 181
18.45. Routable Cellular Modem Interfaces form 181
18.46. Interface Cellmodem HSPA menu 182
18.47. GSM Profile form 182
18.48. Interfaces Cellmodem menu 182
18.49. Cellular Modem Interfaces table 183
18.50. Interfaces Cellmodem HSPA menu 183
18.51. Cellular Modem Interfaces form 183
18.52. HSPA PPP Interfaces Statistics form 184
19.1. S01 Serial Module RJ45 Connector LEDs 185
19.2. Sources of Delay and Error in an End to End Exchange 189
19.3. Serial Protocols menu 192
19.4. Serial Interfaces table 192
19.5. Adding a Protocol in the Edit Private screen 193
19.6. Selecting the Protocol Type in the Edit Private screen 193
19.7. Serial Ports Configuration form 194

19.8. Serial Protocols table 195
19.9. Rawsocket Configuration form 195
19.10. TCP Modbus Configuration form 196
19.11. DNP Protocols Configuration form 198
19.12. DNP Device Address Table Configuration table 198
19.13. DNP Device Address Table Configuration form 198
19.14. Serial Protocol Statistics menu 199
19.15. Serial Port Statistics table 199
19.16. Serial Port Statistics form 200
19.17. Transport Connections Statistics table 201
19.18. TCP/UDP Connection Statistics form 202
19.19. Restart Serserver menu 203
19.20. Restart Serserver Trigger Action 203
19.21. Reset Ports menu 203
19.22. Reset Ports Trigger Action 203
20.1. WAN menu 205
20.2. Interface WAN Slot Port Settings table 205
20.3. Enable WAN Interface form 205
20.4. T1 Parameters form 206
20.5. E1 Parameters form 207
20.6. T3 Parameter form 207
20.7. E3 Parameter form 208
20.8. T1 Channels and Associated Time Slots table 209
20.9. T1 Time Slots form 209
20.10. Adding a Connection 210
20.11. Frame Relay Parameter form 211
20.12. Connection Frame Relay DLCI table 212
20.13. PPP form 212
20.14. Adding an MLPPP Connection 214
20.15. Adding IP and Remote Addresses 215
20.16. T1/E1 Interfaces under the IP submenu 216
20.17. Loopback Test Forms 216
20.18. Loopbacktest Results 216
20.19. WAN Statistics menu 217
20.20. T1E1 Statistics table 217
20.21. Receiving Errors Statistics form 217
20.22. T1E1 Receiving Statistics form 218
20.23. T1E1 Receiving Statistics Form 2 219
20.24. T1E1 Transmitting Errors Statistics form 219
20.25. T1E1 Transmitting Statistics form 220
20.26. T1E1 Transmitting Statistics Form 2 220
20.27. T3 and E3 Alarm form 221
20.28. T1E1 Alarm Indication form 222
20.29. T1E1 Statistics form 223
20.30. PPP Receiving Protocol Statistics form 223
20.31. PPP Transmitting Protocol Statistics form 224
20.32. T1E1 Statistics form 224
20.33. Frame Relay Errors Packets Statistics form 226
20.34. Frame Relay Controlling Packets Statistics form 227
20.35. Frame Relay Receiving Statistics form 228
20.36. Clear Interface Statistics Form And Trigger Action 229
20.37. Clearstatistics Menu Action 229
20.38. Enable Wan Interface form 231
20.39. PPP form 231

20.40. Frame Relay Parameters form 232
20.41. Loopback Test form 233
20.42. DDS Statistics menu 234
20.43. Clear Interface Statistics form 235
21.1. Tunnelling menu 237
21.2. IPsec menu 240
21.3. IPsec form 240
21.4. Syslog form 240
21.5. Show Public RSA Key form 241
21.6. Install-Certificate forms 242
21.7. Install-Ca-Certificate forms 243
21.8. Install-Crl-File forms 244
21.9. Show IPsec Running Status form 244
21.10. Connection table 244
21.11. Connection form 245
21.12. ESP table 246
21.13. ESP Key Settings 246
21.14. IKE table 246
21.15. Public IP Address form 247
21.16. System Public Key form 247
21.17. Nexthop To Other System form 248
21.18. System Identifier form 248
21.19. Private Subnet Behind System form 248
21.20. Network table 249
21.21. Preshared Key table 249
21.22. Preshared Key form 249
21.23. L2TP menu 250
21.24. L2TP form 250
21.25. DNS Server form 250
21.26. PPP Options form 251
21.27. WINS Server form 251
21.28. L2tunneld menu 254
21.29. L2 Tunnel Daemon form 254
21.30. Goose Tunnel table 255
21.31. Goose Tunnel form 255
21.32. Remote Daemon of Goose Tunnel table 255
21.33. Generic L2 Tunnel table 255
21.34. Generic L2 Tunnel Protocol form 256
21.35. Generic L2 Tunnel Egress Interface table 256
21.36. L2 Ethernet Type table 256
21.37. Goose Tunnel Statistics table 256
21.38. Goose Tunnel Statistics form 257
21.39. Connections Statistics table 258
21.40. Connections Statistics form 258
21.41. Generic L2 Tunnel Statistics table 259
21.42. Generic L2 Tunnel Statistics form 259
21.43. Connections Statistics table 260
21.44. Connections Statistics form 260
21.45. Round Trip Time Statistics table 261
21.46. Round Trip Time Statistics form 261
21.47. GRE Example 262
21.48. Generic Routing Encapsulation (GRE) menu 262
21.49. Generic Routing Encapsulation Interfaces table 263
21.50. Generic Routing Encapsulation Interfaces form 263

22.1. OSPF and VRRP Example 269
22.2. Dynamic Routing Menu 270
22.3. RIP Menu 270
22.4. RIP Configuration Form 271
22.5. Routing Timers Form 272
22.6. RIP Interface Parameters Table 274
22.7. RIP Interface Parameters Form 275
22.8. Authentication Form 275
22.9. OSPF Menu 276
22.10. OSPF Configuration Form 277
22.11. OSPF Area Distance Form 278
22.12. Interface Parameters Table 279
22.13. Interface Parameters Form 279
22.14. Dead Interval Form 280
22.15. BGP Menu 281
22.16. BGP Configuration Form 281
22.17. Distance Form 282
23.1. Static Menu 288
23.2. Static Route table 288
23.3. Static Route form 288
23.4. Static Route Using Gateway table 288
23.5. Static Route Using Gateway form 288
23.6. Blackhole Static Route form 289
23.7. Static Route Using Interface table 289
23.8. Static Route Using Interface form 289
24.1. Routing Status Menu 290
24.2. IPv4 Kernel Active Routing Table 290
24.3. IPv6Kernel Active Routing Table 291
24.4. Core Daemon Memory Statistics Form 291
24.5. RIP Daemon Memory Statistics Form 292
24.6. BGP Daemon Memory Statistics Form 292
24.7. OSPF Daemon Memory Statistics Form 293
24.8. RIP Menu 293
24.9. OSPF Menu 293
24.10. Network Table 294
24.11. Reach Table 294
24.12. Router Table 294
24.13. Area Table 295
24.14. Net Table 295
24.15. Summary Table 296
24.16. ASBR-Summary Table 296
24.17. AS-External Table 297
24.18. Neighbor Table 297
24.19. BGP Menu 298
24.20. Route Table 298
24.21. Next Hop Table 299
24.22. BGP Neighbor Table 299
25.1. Multicast Routing menu 301
25.2. Static Multicast Routing Configuration form 301
25.3. Static menu 301
25.4. Multicast Groups Configuration table 301
25.5. Multicast Groups Configuration form 302
25.6. Outgoing Interfaces table 302
25.7. Multicast Routing Status table 303

25.8. Multicast Routing Status form 303
26.1. Security Menu 312
26.2. Firewall Description table 312
26.3. Firewall Description form 312
26.4. Adding a Firewall 313
26.5. Naming a Firewall 313
26.6. Firewall Submenus 314
26.7. Firewall Configuration form 314
26.8. Zone table 315
26.9. Zone form 315
26.10. Main Interface Settings table 316
26.11. Interface Options form 316
26.12. Broadcast Address form 317
26.13. Main Host Settings table 317
26.14. Main Host Settings form 317
26.15. Host Options form 318
26.16. Main Policy Settings table 318
26.17. Main Policy Settings form 318
26.18. Destination Zone form 319
26.19. Source Zone form 319
26.20. Net Address Translation Main Settings table 319
26.21. Net Address Translation Main Settings form 320
26.22. FWMasq table 320
26.23. Net Address Translation Main Settings form 321
26.24. Main Rule Settings table 321
26.25. Main Rule Settings form 322
26.26. Source Zone form 323
26.27. Destination Zone form 323
27.1. Traffic-Control menu 327
27.2. Traffic Control Configuration form 327
27.3. Enabling Basic-configuration Mode 328
27.4. Basic Traffic Control Interfaces table 328
27.5. Interface to Apply Traffic Control form 329
27.6. Basic Traffic Control Priorities table 330
27.7. Priorities form 330
27.8. Enabling Advanced-configuration Mode 332
27.9. Advanced Traffic Control Classes table 333
27.10. TC Classes form 333
27.11. Options form 335
27.12. Advanced Traffic Control Interfaces table 336
27.13. TC Devices form 337
27.14. TCrules menu 338
27.15. Advanced Traffic Control Rules table 338
27.16. TCrules form 339
27.17. Set form 341
27.18. Modify form 342
27.19. Save form 342
27.20. Restore form 342
27.21. Continue form 343
28.1. VRRP Example 345
28.2. VRRP Group Example 346
28.3. VRRP Menu 346
28.4. Virtual Router Redundancy Protocol (VRRP) Form 347
28.5. VRRP Group Table 347

28.6. VRRP Instance Table 347
28.7. VRRP Instance Form 348
28.8. Monitor Interface Form 349
28.9. VRIP IP Address Table 349
28.10. VRRP Status Table 349
28.11. VRRP Status Form 350
29.1. Link Backup Example 351
29.2. Link Fail Over Information Table 352
29.3. Link Fail Over Settings form 353
29.4. Backup Settings form 354
29.5. Link Fail Over Status form 355
29.6. Link Fail Over Logs form 356
29.7. Link Fail Over Test Settings form 357
A.1. The Software Upgrade Menu Interface 359
A.2. Entry Fields in Upgrade Settings Form 360
A.3. Pending Commit 360
A.4. Commit Succeeded 360
A.5. Launch Upgrade 361
A.6. Upgrade Launched Dialogs 361
A.7. Software-Upgrade Menu 361
A.8. Upgrade Monitoring Form in Reboot-pending Stage 362
A.9. Upgrade Monitoring Form Showing Successful Upgrade 363

Preface

This guide describes the web-based user interface for the ROX™ version 2.2 Operating System running on the RuggedRouter® RX1000 family of products.

Supported Platforms

ROX™2.2 is designed to work on RuggedCom's RuggedBackbone™ and RuggedRouter® hardware platforms. This ensures a consistent user experience when migrating from one product model in the family to another.

ROX™ currently supports the following RuggedCom networking platforms:

- RuggedBackbone™ RX5000 family of rugged, modular, Layer 3 switching multi-service hardware platforms.
- RuggedBackbone™ RX1500 family of rugged, modular, hot-swappable Layer 3 switching and routing platforms.
- RuggedRouter® RX1000 family of rugged Cyber-Security Appliances.

Who Should Use This User Guide

This guide is recommended for use by network technical support personnel who are familiar with the operation of networks. Others who might find the book useful are network and system planners, system programmers, and line technicians.

How This Guide Is Organized

Part I, “Administration”

This part covers the graphical user interface and overall management of the hardware chassis and operating system, including access control, logging, networking configuration, and time synchronization.

Part II, “Network Interfaces and Ethernet Bridging ”

This part covers the configuration and monitoring of the Ethernet bridging functions of the system, including Ethernet port setup, the Spanning Tree Protocol, and Virtual LANs.

Part III, “Routing and Security”

This part covers the configuration and monitoring of layer 3 routing and security functions, including OSPF, RIP, BGP, Multicast, and the Firewall.

Each part of this guide is organized into chapters that are typically devoted to one particular feature of the system.

Each chapter discusses mechanisms, protocols, or techniques specific to a particular feature. Many chapters include a general overview of the feature or protocol to be configured, providing some background into the feature and how it is used on the device. All chapters present the forms and fields in the web interface through which you configure the feature.All chapters present the CLI commands you use to configure the feature.

While every effort is made to ensure the accuracy and completeness of this guide, some web interface illustrations may not be exactly as shown.

Applicable Operating System Software Revision

This guide is applicable to ROX™ version 2.2.

Part I. Administration

This part describes the administration of a ROX™-based networking device:

The ROX Web Interface	Chapter 1, *The ROX™ Web Interface*
System Administration	Chapter 2, *System Administration*
Time Synchronization	Chapter 3, *Time Synchronization*
Basic Networking Configuration	Chapter 4, *Basic Network Configuration*
Advanced Networking Configuration	Chapter 5, *IP Network Interfaces*
Alarms	Chapter 6, *Alarms*
Domain Name Search	Chapter 7, *Domain Name Search*
Logging	Chapter 8, *Logging*
SNMP Configuration	Chapter 9, *SNMP*
Authentication	Chapter 10, *Authentication*
Notifications	Chapter 11, *NETCONF*
Physical Chassis Configuration	Chapter 12, *Chassis Management*
PPP User Profiles	Chapter 13, *PPP Users*
DHCP Server	Chapter 14, *DHCP Server*

1. The ROX™ Web Interface

ROX™ features two primary user interfaces: a web-based interface and a command line interface (CLI). This user guide documents the usage and structure of the web-based user interface. For details of the CLI, please refer to the ROX™ Command Line Interface User Guide (in progress).

1.1. Getting Started

1.1.1. Requirements

Accessing the ROX™ web interface for the first time, prior to any system configuration, requires:

- A computer with an installed web browser capable of running JavaScript. ROX™ supports the following web browsers:
  - Microsoft® Internet Exporer 8.0 and higher
  - Mozilla Firefox
  - GNU Iceweasel
  - Google Chrome
- The computer must have a working Ethernet interface, which must be compatible with at least one of the port types on the RuggedRouter® as ordered.
- The ability to configure an IP address and netmask on the computer's Ethernet interface.

1.1.2. Connecting To The Web Interface

By default, the RuggedRouter® RX1000 has a different IP address and subnet configured for each of four distinct IP interfaces:

Interface Name	IP Address/Mask
fe-3-1	192.168.0.1/24
fe-3-2	192.168.1.1/24
fe-4-1	192.168.2.1/24
fe-4-2	192.168.3.1/24

Table 1.1. Default IP Address Configuration

In order to connect to the RX1000 using a web browser, configure the IP address of the web browser's system to fall within the subnet of the corresponding RX1000 interface. For example, if the web browser system is connected to the fe-3-2 Ethernet interface:

- The web browser system's Ethernet interface must be configured with an IP address in the range: 192.168.1.2 to 192.168.1.254.
- The RX1000 is accessible to the web browser at the IP address: 192.168.1.1, the address of the fe-3-2 network interface.

1.1.3. The Web Browser Connection

The ROX™ web server uses SSL (Secure Socket Layer) to encrypt data traffic exchanged with its clients (connections made via "https://"). This guarantees the privacy of communications between browser and server.

It can happen that upon connecting to the ROX™ web server, some new web browsers may report that they cannot verify the authenticity of the server's certificate against any of their known certificate authorities. This is expected, and it is safe to instruct the browser

to accept the certificate offered by the ROX™ system. Once the browser is instructed to accept the certificate, all communications with the web server will be secure.

Start a web browser session and open a connection to the switch by entering a URL that specifies its IP address (https://192.168.1.2, to continue with the example above). Once the web browser makes contact with the switch, The resulting page should be the login prompt displayed below:

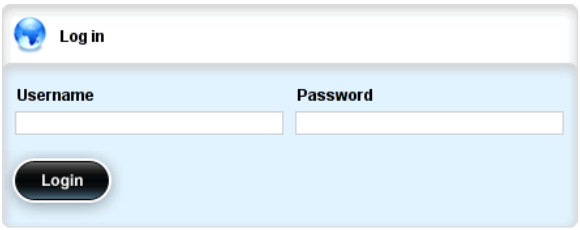

Figure 1.1. The ROX™ Login Form

Enter the default user name, "admin" and the configured password for the admin user. Click on the **Login** button. The switch is shipped with a default administrator password, "admin". If authentication is successful, the main menu is presented.

1.2. The Structure Of The Web Interface

The system configuration interface (the *Configure Running* tab) is organized as a hierarchical set of linked menu entries, which may be traversed using the four-panel navigation window, as illustrated below.

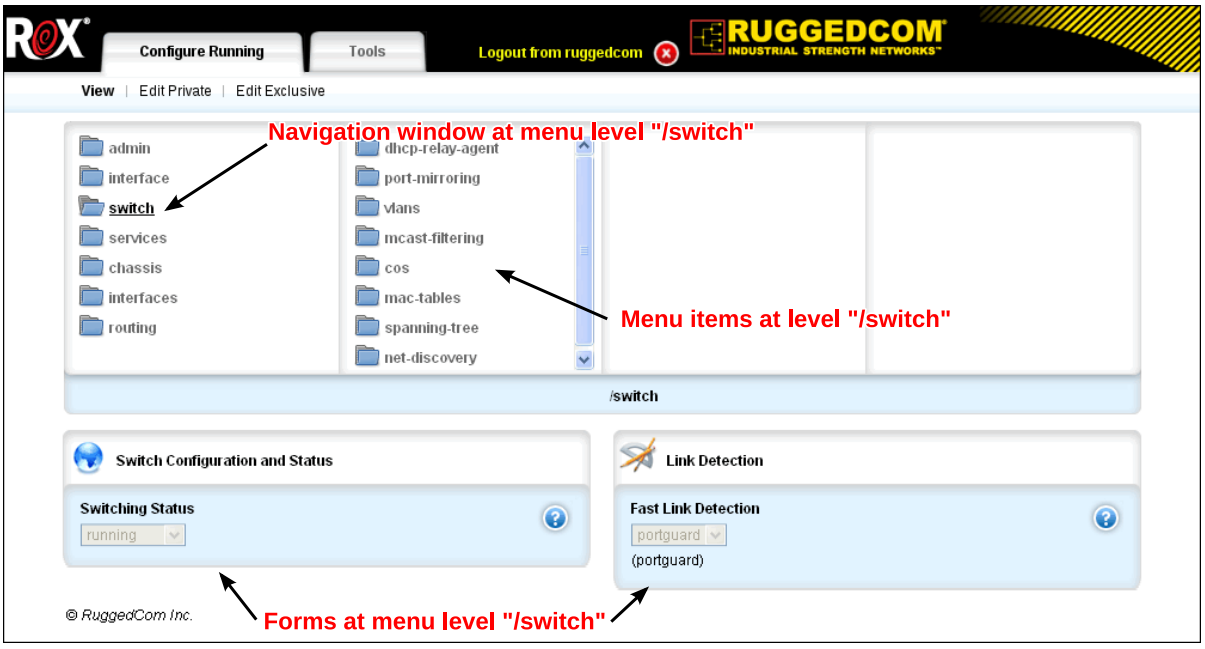

Figure 1.2. The ROX™ Web Interface

Menu items listed in a panel of the navigation window at a given point in the menu hierarchy may be:

- Submenus, which are marked using the icon, or
- Actions, which are marked using the icon.

Note that green submenu icons represent operational data.

Tables and forms relevant to the selected menu item appear below the navigation window.

The icons in the upper left corner of the forms and tables are used to signify the type of content represented in each form or table.

- The green arrow icon signifies operational data.
- The key icon signifies the key in key settings.
- The blue globe icon signifies the global group (a high-level grouping of items).
- The pencils and protractor icon signifies configuration data.
- The paper and pencil icon signifies results. This icon is usually found on a form where there are parameters to enter.

Every web page in the ROX™ user interface has a header, illustrated above, containing:

- The ROX™ and RuggedCom logos and a **Logout** button, which terminates the current web session.
- The tabs: *Configure Running* and *Tools*.

The *Configure Running* tab selects the configuration interface described above. A menu bar below the page header displays the following editing mode controls:

- *View*: View configuration settings only.
- *Edit Private*: Enter a configuration editing mode where you can make changes to the system. Your changes are applied to the active system only when you commit them. Edit sessions are self contained: the changes made in your edit session are not visible to other users in other edit sessions.
- *Edit Exclusive*: Enter a configuration editing mode where, after committing your changes, you can specify a timeout period to test the changes. At the end of the timeout period, your changes to revert back to the original settings. Use this mode when you want to test changes before committing them permanently. When you click **Commit**, a dialog prompts you to set a commit timeout. Type a value and select a unit of time. ROX temporarily applies your changes to the active system for the specified time. To cancel the commit and discard the changes, click **Abort Commit** before the time elapses. To permanently commit your changes, click **Commit** before the time elapses.

In many cases, the tables appear on a screen closer to the top level and clicking on one of the submenus brings up the form(s) associated with the table. For example, clicking on the Chassis menu and then the Hardware submenu will display the Slot Hardware table. Further clicking on the pm1 submenu will display the Slot Hardware form.

The *Tools* tab displays a menu of tools in the menu bar, with the following structure:

- Device Info: displays text from various system logs. You can specify the number of lines to view and a text filter.
  - Messages Viewer: displays all events from **`/var/log/messages`**.
  - Syslog Viewer: displays syslog events from **`/var/log/syslog`**.
  - Authlog Viewer: displays authentication events from **`/var/log/auth.log`**.
  - Layer2log Viewer: displays Layer 2 events from **`/var/log/layer2`**.
  - Kernlog Viewer: displays kernel events from **`/var/log/messages`**.
- Accessories
  - Ping: an ICMP echo tool for IPv4 addresses
  - Ping6: an ICMP echo tool for IPv6 addresses

- Tcpdump: a packet analyzer for TCP/IP and other packets
- Traceroute: a tool for displaying route or path information and packet transit delays between IPv4 addresses
- Traceroute6: a tool for displaying route or path information and packet transit delays between IPv6 addresses

- CLI: a command line interface window
- Users: displays a list of currently connected users, provides controls to kick users off of the system, and provides a message board to send messages to users.
- Upload: uploads configuration files, feature keys, elan certificates, ipsec certificates, ca certificates, and crl certificates to the system from your workstation. From the *Choose file type* list, select the type of file to upload. Click **Choose File** and navigate to and select a file on your workstation. To upload the selected file, click **Send**.
- Download: downloads configuration files, feature keys, elan certificates, ipsec certificates, ca certificates, crl certificates, log files, and rollback files from the system to your workstation. From the *Choose file type* list, select the type of file to download. Click **List files**; a list of available files appears. To download a file, right-click on a file name and select **Save Link As** (the name of the menu option will vary, depending on your browser). To open a file in a new window or tab, click on a file name.

1.2.1. Top-level Menu Categories

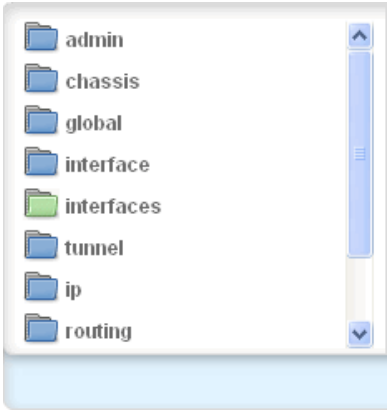

Figure 1.3. Top-level Menu

Below is a description of the categories in the top-level menu that is shown above.

admin

The admin menu is used for configuring functions related to the administration of the router. Functions include DNS, alarms, logging, authentication, users, software upgrade, notifications and SNMP.

chassis

The chassis menu is used for configuring the chassis.

global

The global menu is used for configuring global functions including profiles for PPP and cellular modems.

interface

The interface menu is used for configuring the interface, including (where applicable) sections for WAN, serial, modem and trunks.

interfaces

The interfaces menu displays the status of functions configured via the interface menu. For example, eth functions can be configured using the eth submenu that is accessible from the interface menu. The eth status can be viewed by clicking on the eth submenu of the interfaces menu.

tunnel

The tunnel menu is used for configuring IP tunnels using IPsec, Layer 2 tunnelling functions and Generic Routing Encapsulation (GRE).

ip

The ip menu is used for configuring the ROX™ system's IP network interfaces.

qos

The qos menu is used for configuring traffic control.

routing

The routing menu is used for configuring the routing features. Included are sections on dynamic, static, status and multicast routing.

security

The security menu is used for configuring security, including the firewall.

services

The services menu is used for configuring various services. These services include timekeeping, VRRP, DHCP server and linkfailover.

1.3. Making Configuration Changes

In order to make configuration changes, select the desired *Edit Private* mode from the configuration view. The same navigation window, tables and forms are redisplayed, but with additional controls, as illustrated below.

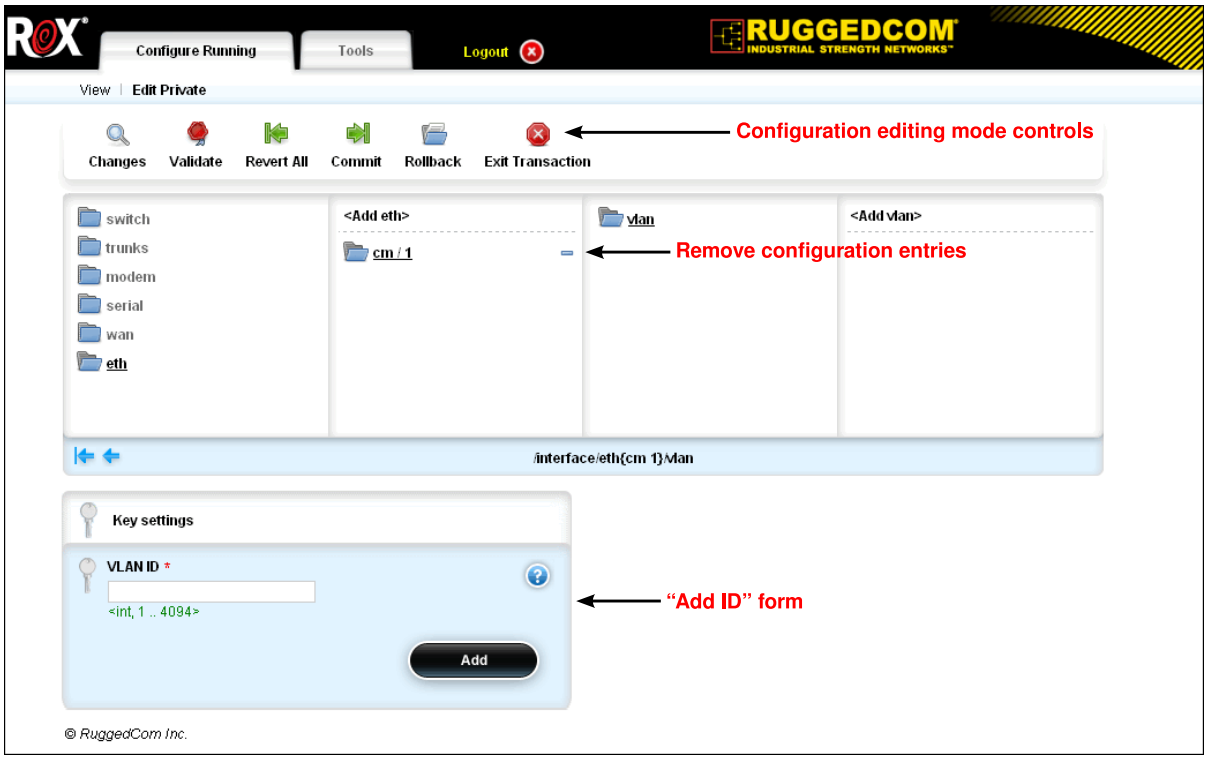

Figure 1.4. Example of Edit Private Mode

The example above depicts the process of adding a VLAN ID to an interface. The **interface/eth/cm1** menu can be seen to contain:

- A configuration entry, followed by a "delete" icon, , which removes the corresponding entry.

Clicking on <add vlan> displays the *Add ID* form below the navigation window, which prompts for a VLAN ID. Entering a VLAN ID and clicking **Add** adds the selected VLAN to the currently selected interface.

Note the help button, , on the *Add ID* form which, when clicked, displays context-sensitive information about the corresponding data field.

A red asterisk appears beside fields that are mandatory for configuration, when in Edit Private mode. Note the red asterisk next to the field name (VLAN ID) in the Key settings form.

Several controls below the header and menu bar are used to affect the behaviour of the changes made during the current configuration editing session:

Changes

Present a summary of all pending changes.

Validate

Automatically check the validity of pending changes.

Revert All

Abort all pending changes.

Commit

Commit all pending changes - save changes persistent configuration storage and to the running system.

Rollback

Present a list of change sets made to date, with an option to revert a selected set of changes.

Exit Transaction

Exit from configuration editing mode. If there are pending changes, a prompt will be presented to verify the discarding of all pending changes.

1.3.1. Configuring Tables Using Key Settings Forms

Much of the information in ROX™ is organized into tables. Each table is indexed or sorted by a key, which is a piece of information such as a name, address, or other variable. For example, a Chassis Hardware table is indexed by slot name (with the slot name being the key) and a DNS Server table is indexed by IP address (with the IP address being the key). Key information can be added using the key settings forms. To add server information to a DNS server table, for example, add the server address to the key settings form and this information will appear in the DNS server table.

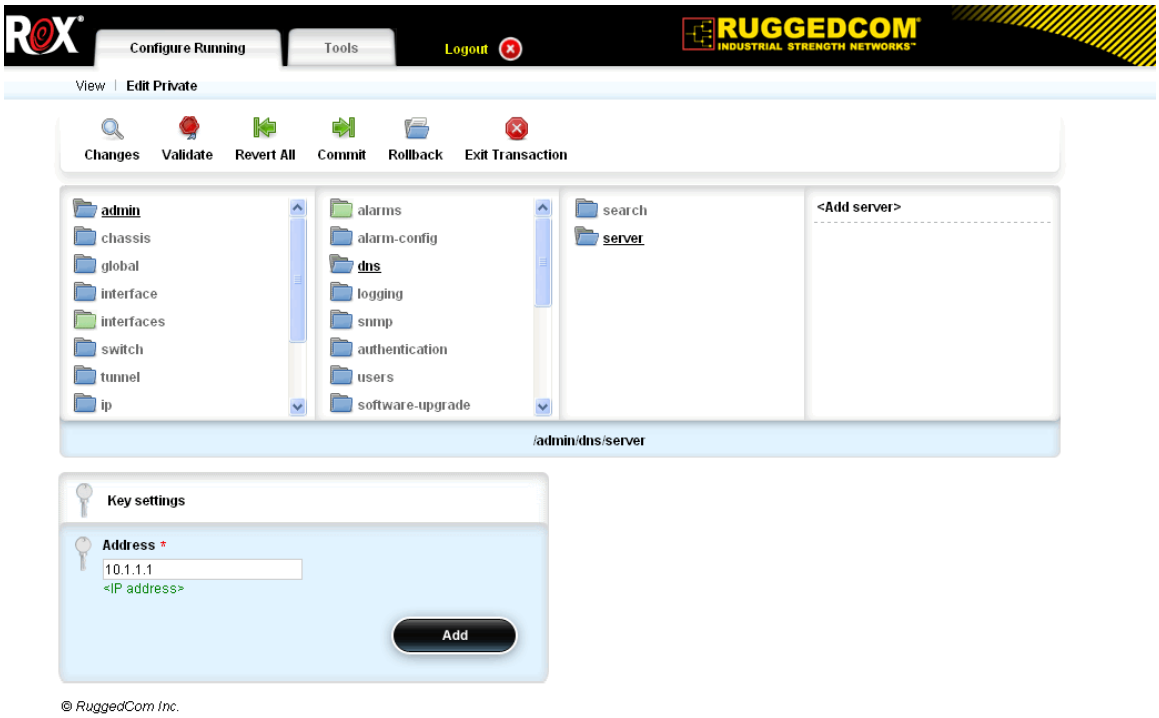

Figure 1.5. Adding Key Information

To add key information to a table, go into the *Edit Private* mode and enter the information into the key settings form. Click the *Commit* button. When you have finished making all changes, click the *Exit Transaction* button to return to the *View* mode.

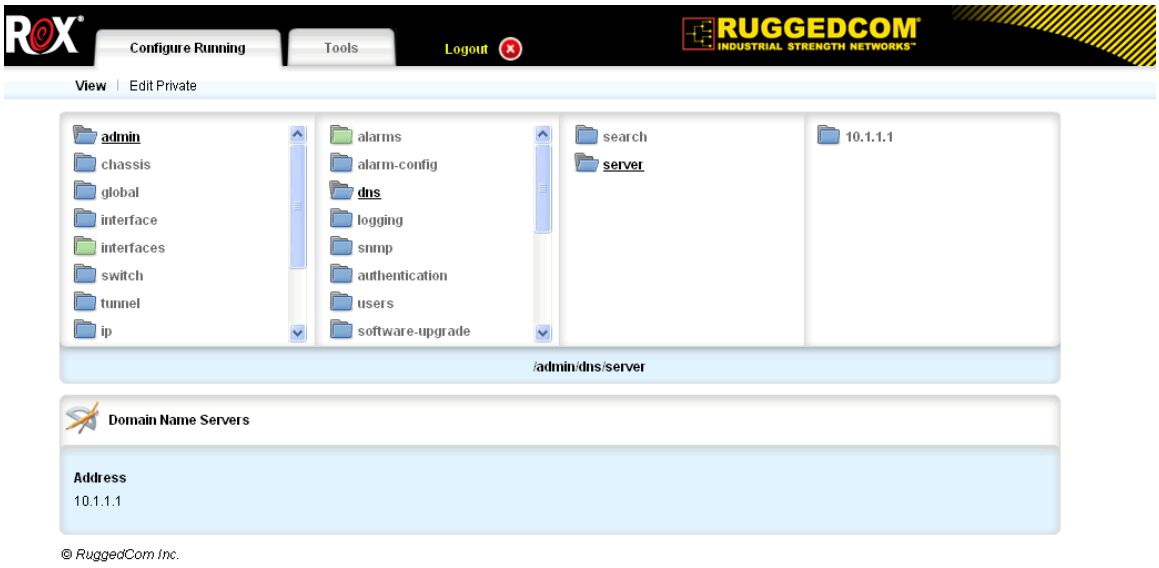

Figure 1.6. Key Information in a Table

The information entered in the key settings form will now appear in the table. Note that the table appears on the server screen, while the key settings form appears on the address screen, which is a submenu linked to the server screen (see below).

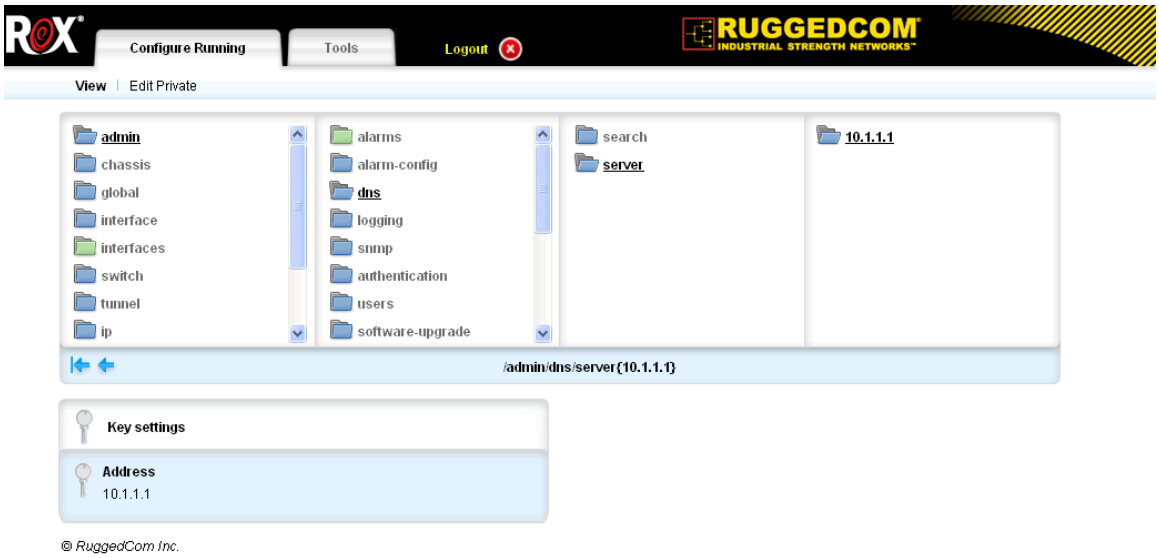

Figure 1.7. Example of Key Settings 1

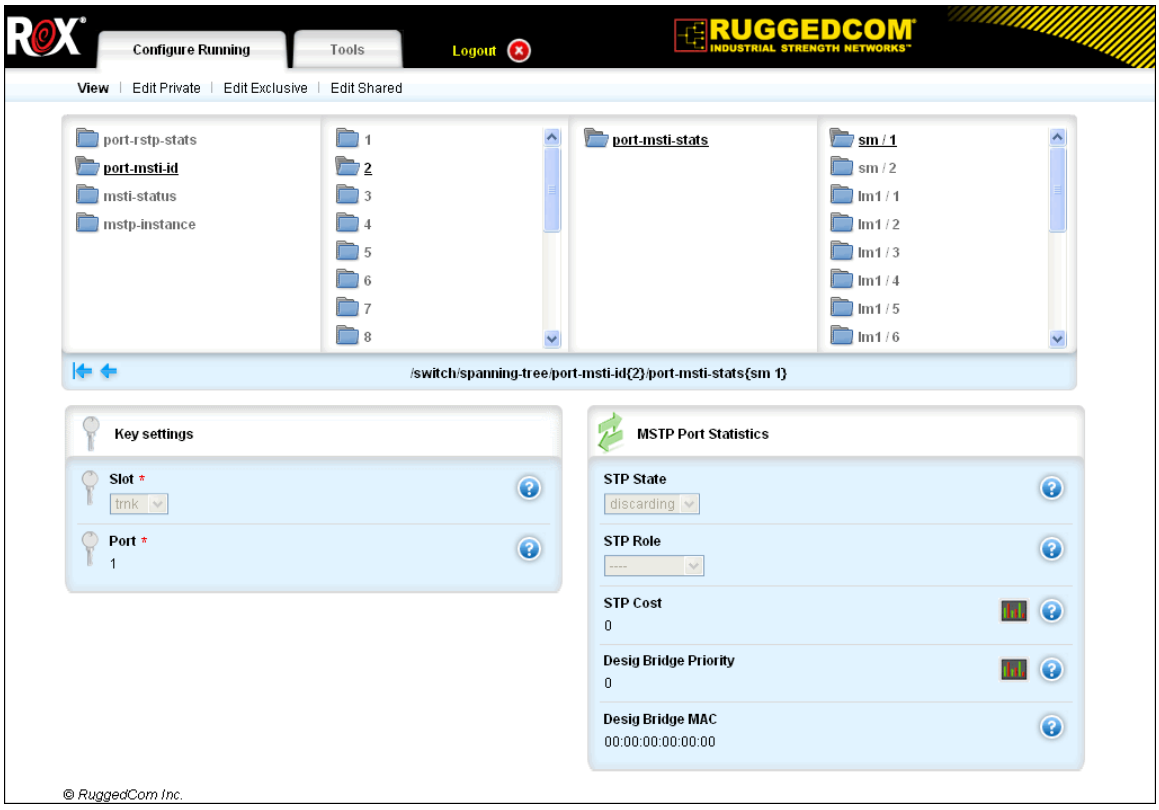

Figure 1.8. Example of Key Settings 2

The submenus that display the key settings forms appear in the far right column of the screen. Sometimes, it will be necessary to traverse several menu screens to get to a key settings form.

1.3.2. Viewing More Information in Tables

Occasionally, a table may have more entries that are not visible in the initial view. If you encounter a table that has a line of linked text at the top with the word "Next", and a number in parentheses ( ), you

can click on the "Next" link and access additional entries. The two figures below illustrate this situation. In this case, there are 18 entries in the table. The first table contains 16 entries and 2 entries follow in the next table.

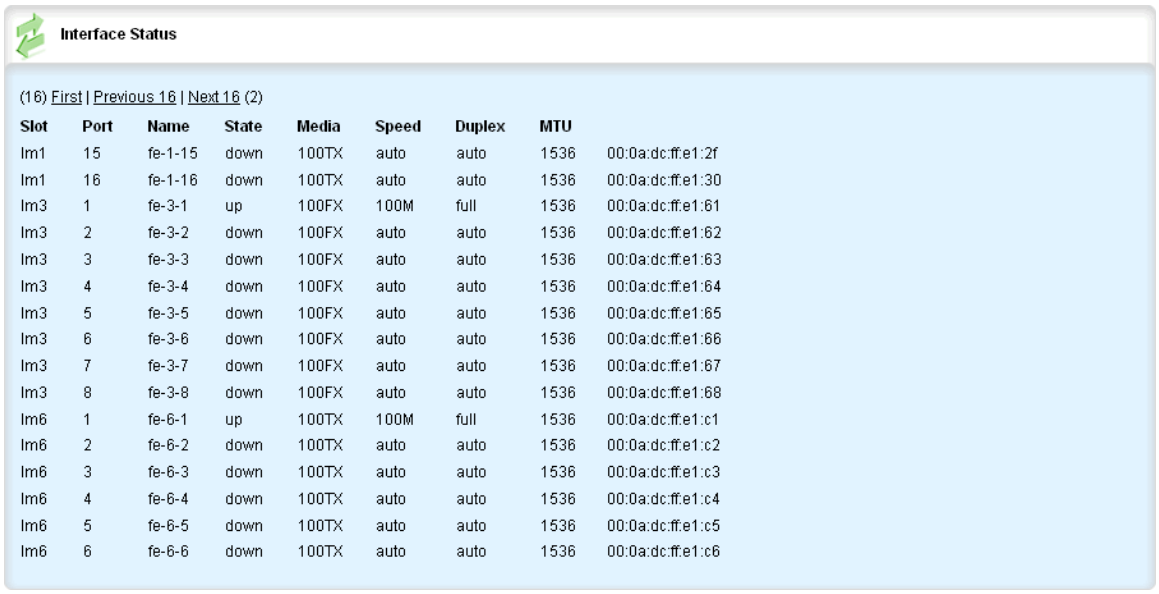

Interface Status

(16) First | Previous 16 | Next 16 (2)

Slot	Port	Name	State	Media	Speed	Duplex	MTU	
lm1	15	fe-1-15	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:2f
lm1	16	fe-1-16	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:30
lm3	1	fe-3-1	up	100FX	100M	full	1536	00:0a:dc:ff:e1:61
lm3	2	fe-3-2	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:62
lm3	3	fe-3-3	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:63
lm3	4	fe-3-4	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:64
lm3	5	fe-3-5	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:65
lm3	6	fe-3-6	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:66
lm3	7	fe-3-7	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:67
lm3	8	fe-3-8	down	100FX	auto	auto	1536	00:0a:dc:ff:e1:68
lm6	1	fe-6-1	up	100TX	100M	full	1536	00:0a:dc:ff:e1:c1
lm6	2	fe-6-2	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c2
lm6	3	fe-6-3	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c3
lm6	4	fe-6-4	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c4
lm6	5	fe-6-5	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c5
lm6	6	fe-6-6	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c6

Figure 1.9. First Table of Information

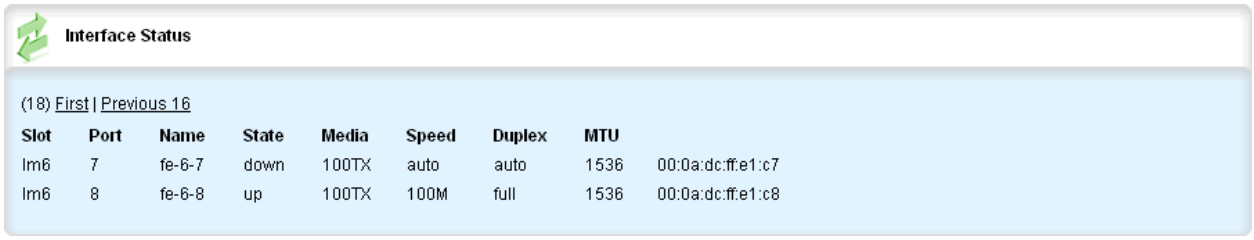

Interface Status

(18) First | Previous 16

Slot	Port	Name	State	Media	Speed	Duplex	MTU	
lm6	7	fe-6-7	down	100TX	auto	auto	1536	00:0a:dc:ff:e1:c7
lm6	8	fe-6-8	up	100TX	100M	full	1536	00:0a:dc:ff:e1:c8

Figure 1.10. Second Table of Information

The second table of information shows the balance of the entries and contains a link back to the previous entries.

2. System Administration

This chapter describes administration-related functions and the Administration menu. Information on the Administration submenus is found throughout Part 1 of this guide.

2.1. Administration menu

Figure 2.1. Administration menu

The Administration (Admin) menu is accessible from the main menu. Use this menu to link to submenus related to alarms, DNS, logging, SNMP, authentication, user IDs and passwords, software versions (upgraded) and netconf.

As well, you can link directly from the Admin menu to commands called "actions" (see below) that will clear or acknowledge all alarms, shut down or reboot the system, set the system clock or restore factory defaults.

2.2. System Commands

This section describes where to find basic system commands using the Administration menu and its menu actions. The following forms are accessible from the Administration menu.

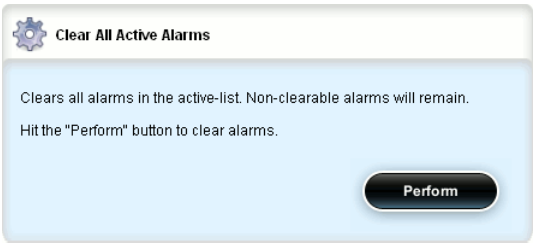
Clear All Active Alarms
Clears all alarms in the active-list. Non-clearable alarms will remain.
Hit the "Perform" button to clear alarms.
Perform

Figure 2.2. Clear All Alarms Menu Action form

To clear all alarms, click on the *clear-all-alarms* menu action and then click the *Perform* button on the Clear All Alarms form.

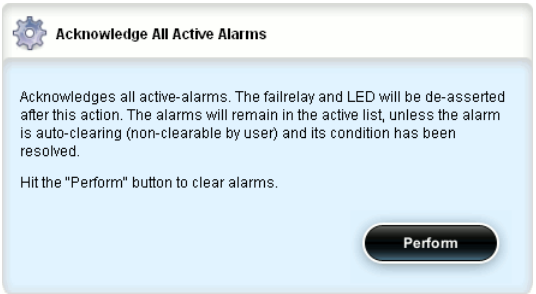
Acknowledge All Active Alarms
Acknowledges all active-alarms. The failrelay and LED will be de-asserted after this action. The alarms will remain in the active list, unless the alarm is auto-clearing (non-clearable by user) and its condition has been resolved.
Hit the "Perform" button to clear alarms.
Perform

Figure 2.3. Acknowledge All Alarms Menu Action form

To acknowledge all alarms, click on the *acknowledge-all-alarms* menu action and then click the *Perform* button on the Acknowledge All Alarms form.

Figure 2.4. Shutdown the Device Menu Action form

To shut down the device, click on the *shutdown* menu action and then click the *Perform* button on the Shutdown the Device form.

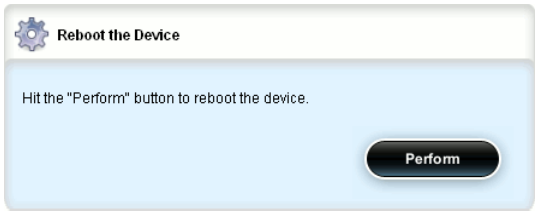

Figure 2.5. Reboot the Device Menu Action form

To reboot the device, click on the *reboot* menu action and then click the *Perform* button on the Reboot the Device form.

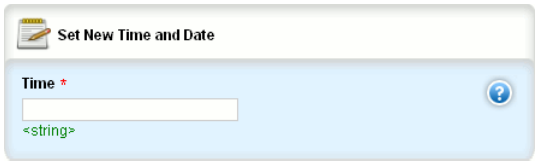

Figure 2.6. Set New Time and Date form

The Set New Time and Date form configures the current time and date settings.

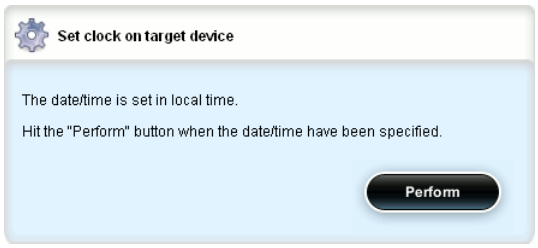
Set clock on target device
The date/time is set in local time.
Hit the "Perform" button when the date/time have been specified.
Perform

Figure 2.7. Set Clock on Target Device form

To set the clock on the target device, click on the *setSystemClock* menu action, then enter the relevant time/date information into the Set New Time and Date form. The information must be in the following format: YYYY-MM-DD HH:MM:SS. After entering this information, click the *Perform* button on the Set clock on target device form.

For more detailed information on time synchronization, refer to Chapter 3, *Time Synchronization*.

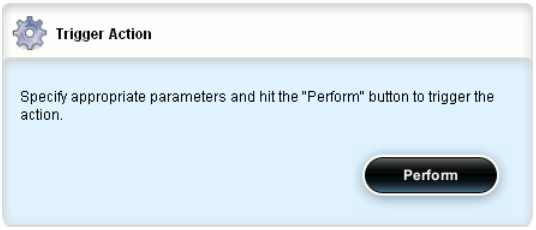

Figure 2.8. Restore-factory-defaults Trigger Action form

To restore factory defaults to the system, click on the *restore-factory-defaults* menu action and then click the *Perform* button on the Restore-factory-defaults Trigger Action form.

The Administration, Hostname, Timezone and Current System Time forms are accessible from the Admin menu.

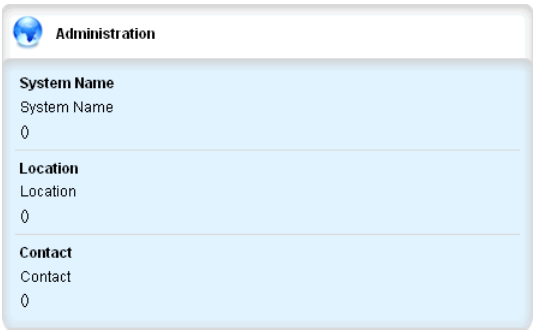

Figure 2.9. Administration form

System Name

Synopsis: *A string*

Default: *System Name*

An administratively-assigned name for this managed node. By convention, this is the node's fully-qualified domain name. If the name is unknown, the value is the zero-length string.

Location

Synopsis: *A string*

Default: *Location*

The physical location of this node (e.g., 'telephone closet, 3rd floor'). If the location is unknown, the value is the zero-length string.

contact

Synopsis: *A string*

Default: *Contact*

The textual identification of the contact person for this managed node, together with information on how to contact this person. If no contact information is known, the value is the zero-length string.

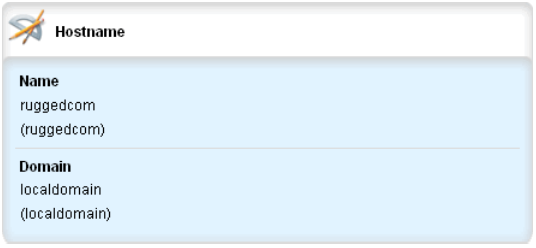

Figure 2.10. Hostname form

The hostname is the name of the product. (This can be changed, though.)

name

Synopsis: *A string conforming to: "[A-Za-z0-9]([A-Za-z0-9-]*[A-Za-z0-9])*"*

Default: *ruggedcom*

The hostname is the name of this device.

domain

Synopsis: *Domain name (RFC 1034)*

Default: *localdomain*

The domain for this hostname.

Figure 2.11. Timezone form

Timezone Category

Synopsis: *string*

Selects the timezone. Note that the Etc/GMT timezones conform to the POSIX style and have their signs reversed from common usage. In POSIX style, zones west of GMT have a positive sign; zones east of GMT have a negative sign.

Timezone

Synopsis: *string*

Selects the timezone.

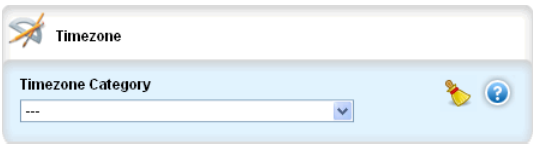

Figure 2.12. Setting the Timezone Form - in Edit Private Mode

To set the time zone, enter Edit Private mode and click on the Timezone Category field. Use the drop-down menu which appears to select the appropriate time zone. Daylight saving time will adjust automatically, if applicable to your zone.

Figure 2.13. Current System Time form

The Current System Time form displays the current time.

UTC Time

Synopsis: *string*

The current GM Time

Local Time

Synopsis: *string*

The current local time

2.3. Administrative Access Control

The following access control forms are accessible from the Administration menu - by clicking on the main menu under *admin*.

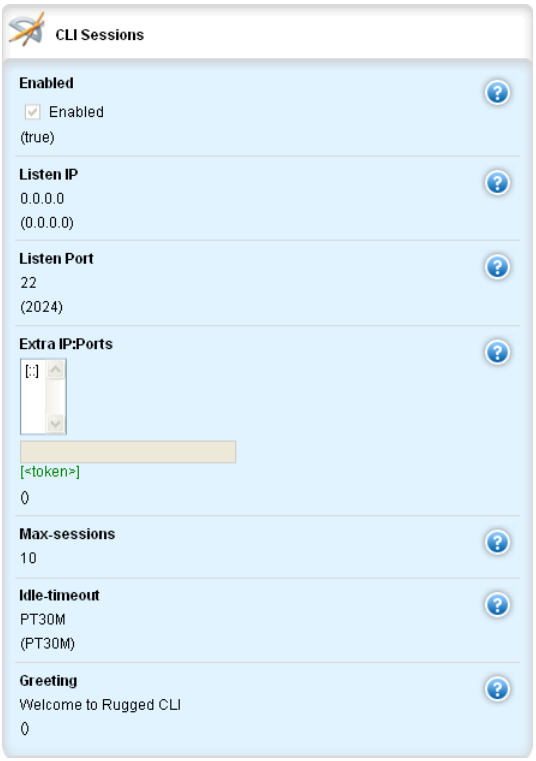

Figure 2.14. CLI Sessions form

enabled

Synopsis: *boolean*

Default: *true*

Provides the ability to configure CLI features on the device.

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the CLI will listen on for CLI requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *22*

The port on which the CLI listens for CLI requests. The default is port 22.

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

The CLI will also listen on these IP Addresses:Port values. Add ':#' to set non-default port value #. (ie. xxx.xxx.xxx.xxx:19343 [::] [::]:16000)

Maximum Number of CLI Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *10*

The maximum number of concurrent CLI sessions

Idle Timeout

Default: *PT30M*

Maximum idle time before terminating a NETCONF session. If the session is waiting for notifications, or has a pending confirmed commit, the idle timeout is not used. The default value is 0, which means no timeout.

greeting

Synopsis: *string*

Default: *Welcome to Rugged CLI*

Sets the greeting presented when the user logs in to the CLI.

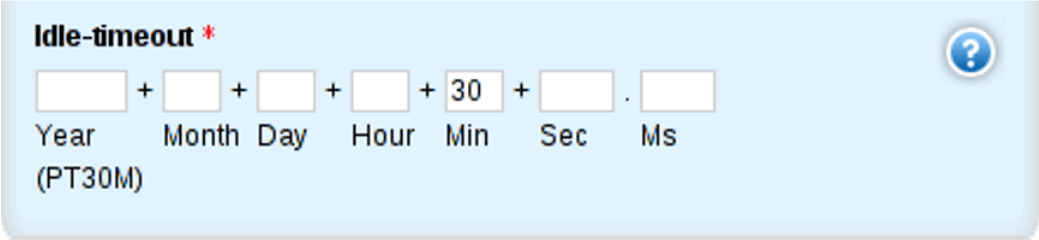

Figure 2.15. Idle-timeout field

Clicking on the Idle-timeout field on the CLI Sessions form allows you to choose a value for this field. The default value is PT30M, which stands for "Precision Time 30 Minutes". This refers to the time when an inactive session expires or times out. Only integer values corresponding to the following fields can be entered: Year, Month, Day, Hour, Min, Sec, or Ms. The example above shows the default value of PT30M, which corresponds to the Min field.

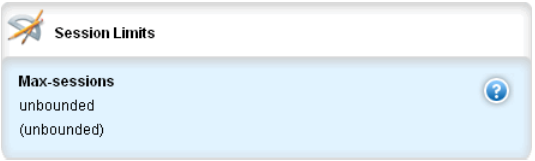

Figure 2.16. Session Limits form

The Session Limits form is used for setting the maximum number of users sessions on a northbound channel.

Maximum Sessions Total

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *70*

Puts a limit to the total number of concurrent sessions to ROX 2.

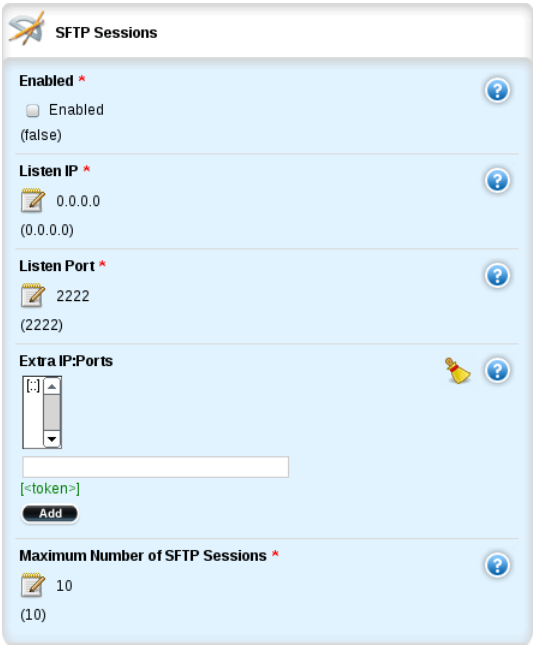

Figure 2.17. STFP Sessions form

The SFTP Sessions form sets the parameters for Secure File Transfer Protocol (SFTP) sessions.

enabled

Synopsis: *boolean*

Default: *false*

Enable/Disable the SFTP user interface

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the SFTP will listen on for SFTP requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *2222*

The port the SFTP will listen on for SFTP requests (default 2222).

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

The SFTP will also listen on these IP Addresses:Port values. Add ':#' to set non-default port value #. (ie. xxx.xxx.xxx.xxx:19343 [::] [::]:16000)

Maximum Number of SFTP Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *10*

The maximum number of concurrent SFTP sessions

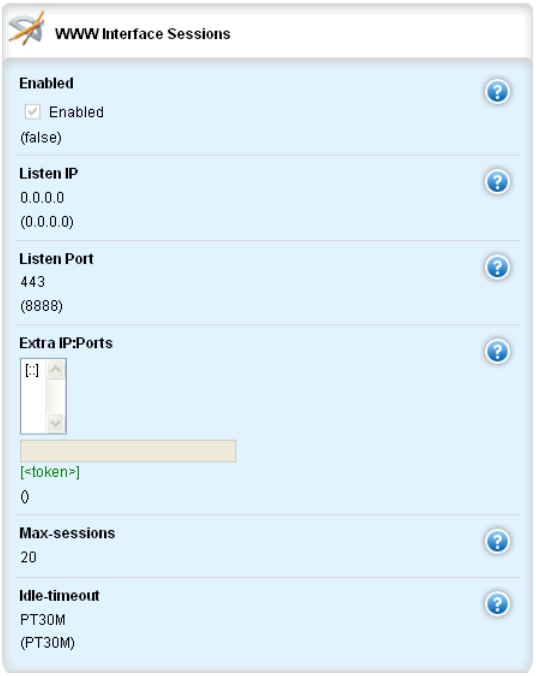

Figure 2.18. WWW Interface Sessions

The WWW Interface Sessions form provides control of WWW User Interface settings.

enabled

Synopsis: *boolean*

Default: *true*

Provides the ability to configure WebUI features on the device.

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the CLI will listen on for WebUI requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *443*

The port on which the WebUI listens for WebUI requests. The default is port 443.

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

The WebUI will also listen on these IP Addresses:Port values. Add ':#' to set non-default port value #. (ie. xxx.xxx.xxx.xxx:19343 [::] [::]:16000)

Maximum Number of WebUI Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *20*

The maximum number of concurrent WebUI sessions

Idle Timeout

Default: *PT30M*

Maximum idle time before terminating a WebUI session. If the session is waiting for notifications, or has a pending confirmed commit, the idle timeout is not used. The default value is 0, which means no timeout.

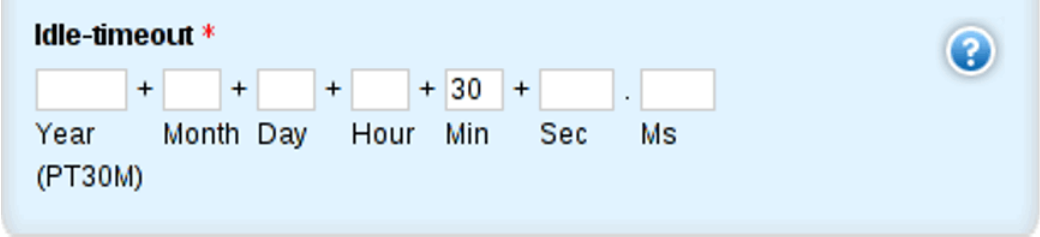

Figure 2.19. Idle-timeout field

Clicking on the Idle-timeout field on the WWW Interface Sessions form allows you to choose a value for this field. The default value is PT30M, which stands for "Precision Time 30 Minutes". This refers to the time when an inactive session expires or times out. Only integer values corresponding to the following fields can be entered: Year, Month, Day, Hour, Min, Sec, or Ms. The example above shows the default value of PT30M, which corresponds to the Min field.

2.4. User Accounts

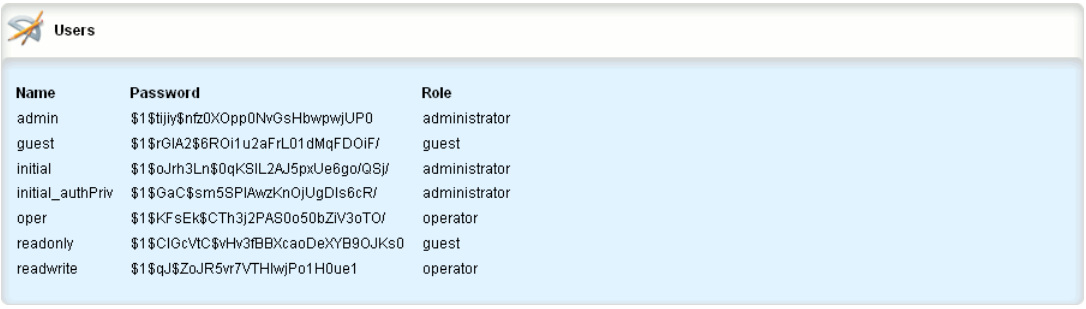

Figure 2.20. Users menu

The Users menu is accessible from the main menu under *admin*. This menu is used to access commands needed for creating and managing passwords for administrators, operators and guests. Both private and public passwords can be created. The Admin Users ID Table (below) can be found on the same screen as the Users menu. Clicking on admin, guest, oper, private or public will lead you to the Users ID forms for each of these options.

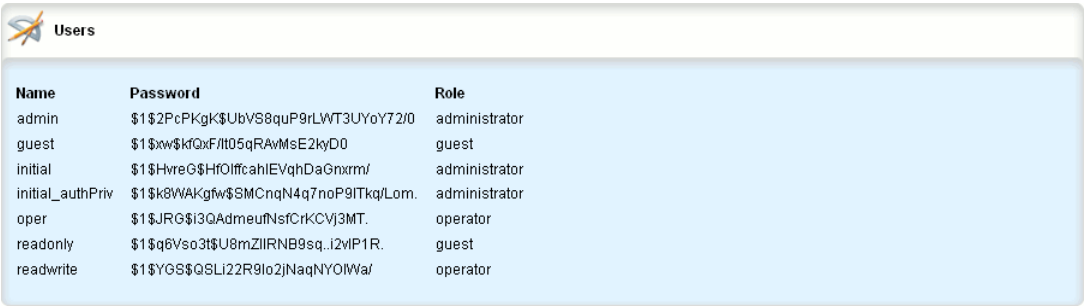

Figure 2.21. Users table

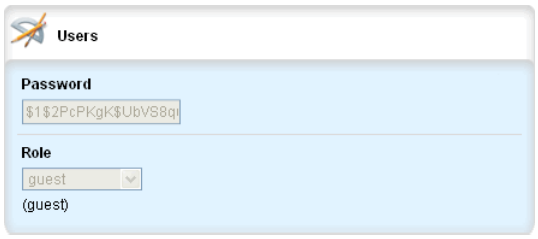

Figure 2.22. Users form

name

Synopsis: *string*

User Name

password

Synopsis: *A string*

User Password

role

Synopsis: *string - one of the following keywords { guest, operator, administrator }*

Default: *guest*

User Role

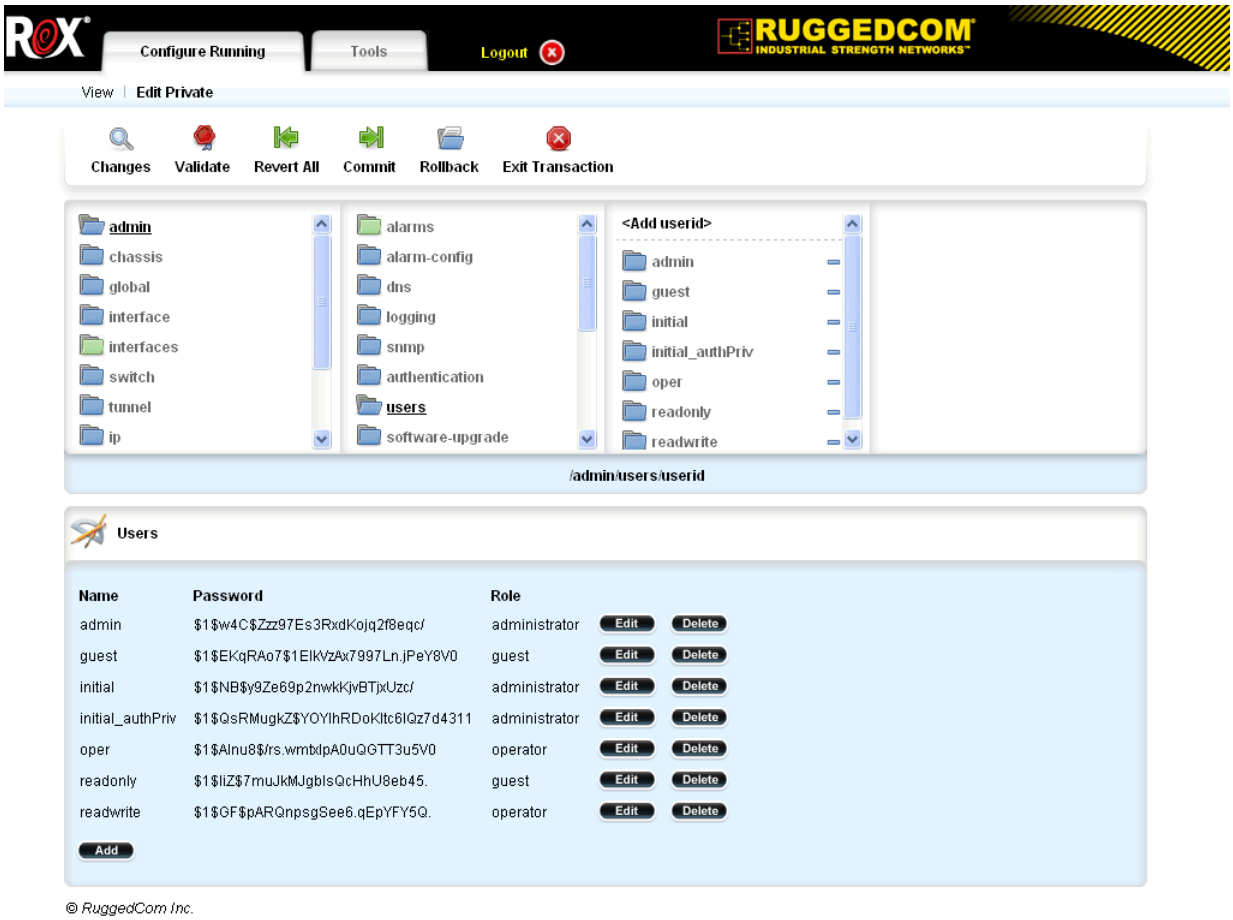

Figure 2.23. Users Screen in Edit Private View

Passwords can be managed, added and deleted while in the Edit Private view.

2.5. Software Upgrade

ROX™ supports two system partitions. One is always active and the other is inactive. ROX™ always applies software upgrades to the inactive partition, providing the following advantages:

1. The current system is unaffected and can operate normally while the upgrade is in progress
2. The current partition remains intact, allowing you to roll back to the original system if needed

After a successful upgrade, the next reboot boots the upgraded partition.

The following applies to software upgrades:

- All system configurations and all user files (featurekeys, configuration files etc.) are carried over to the upgraded partition.
- All configurations are locked during an upgrade and until the upgraded partition is booted. This prevents post-upgrade configuration changes that are not carried over to the upgraded partition.
- Completed upgrades can be declined before the next reboot.
- If major system failures are detected upon booting the upgraded partition, the system will automatically roll back to the previous partition.

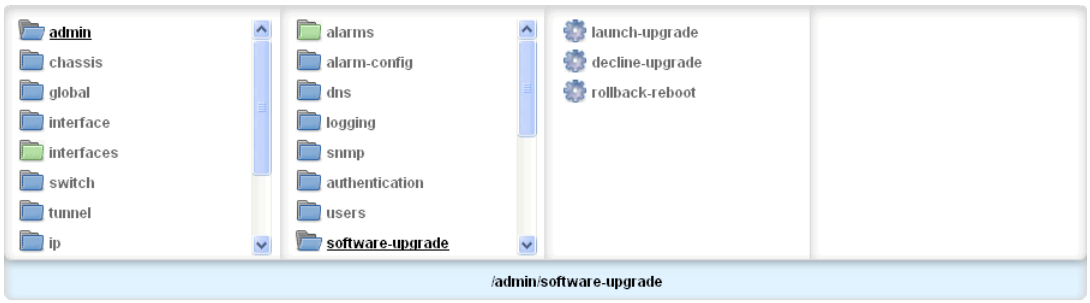

Figure 2.24. Software-Upgrade menu

The Software-Upgrade menu is accessible from the main menu under *admin*. The path to this menu is *admin/software-upgrade*. This menu links to functions that will enable the user to upgrade software, launch the upgraded software, decline new upgrades, and rollback and reboot. The Upgrade Monitoring form and Upgrade Settings form appear on the same screen as the Software-Upgrade menu.

Figure 2.25. Upgrade Settings

In edit mode, define an upgrade server on the *Upgrade Settings* form by setting the Server URL and Target ROX Version parameters. The Upgrade Server URL is the location of the ROX™ software repository. Target ROX Version is the version of ROX to which you are upgrading. For information on setting up an upgrade server, see Appendix C, *Setting Up An Upgrade Server*.

Upgrade Server URL

Synopsis: *string*

repository-url

Target ROX Version

Synopsis: *string*

target-version

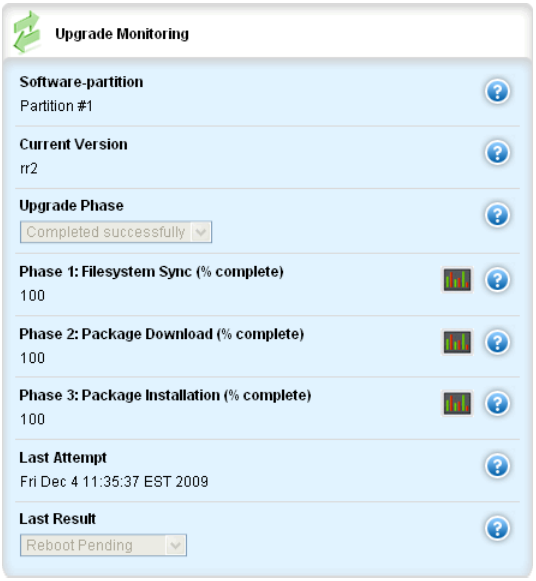

Figure 2.26. Upgrade Monitoring

The Upgrade Monitoring form displays the status of the current upgrade operation.

software-partition

Synopsis: *A string*

The current active partition number. The unit has two software partitions: #1 and #2. Upgrades are always peformed to the other partition.

Current Version

Synopsis: *A string*

The current operating software version.

Upgrade Phase

Synopsis: *string - one of the following keywords { Failed, Completed successfully, Unknown state, Installing packages, Downloading packages, Copying filesystem, Estimating upgrade size, Inactive }*

The current phase or state of the upgrade. It is one of 'Estimating upgrade size', 'Copying filesystem', 'Downloading packages', 'Installing packages', Unknown state', 'Completed successfully', or 'Failed'. These phrases will not vary, any may be used programmitcally for ascertaining state.

status-message

Synopsis: *string*

Additional details on the status of the upgrade

Phase 1: Filesystem Sync (% complete)

Synopsis: *integer*

Phase 1 of the upgrade involves synchronizing the filesystem with the partition to which we are upgrading.

This reflects the estimated percent complete.

Phase 2: Package Download (% complete)

Synopsis: *integer*

Phase 2 of the upgrade downloads all packages that require an update. This reflects the estimated percent complete.

Phase 3: Package Installation (% complete)

Synopsis: *integer*

Phase 3 of the upgrade installs all packages that require an update. This reflects the estimated percent complete.

Last Attempt

Synopsis: *A string*

The date and time of completion of the last upgrade attempt.

Last Result

Synopsis: *string - one of the following keywords { Interrupted, Declined, Not Applicable, Reboot Pending, Unknown, Upgrade Failed, Upgrade Successful }*

Indicates whether or not the last upgrade completed successfully

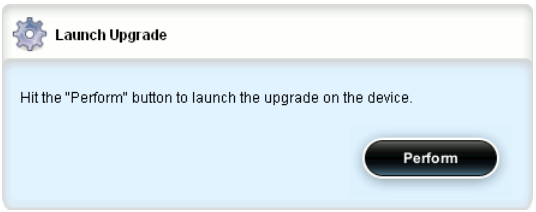

Figure 2.27. Launch Upgrade

To launch an upgrade, click on the *launch-upgrade* menu action and then click the *Perform* button on the Launch Upgrade form. Note that the server URL and version name information must be entered in the Upgrade Settings form prior to launching the upgrade. For detailed step-by-step instructions on how to perform a software upgrade, refer to Appendix A, *Upgrading Software*.

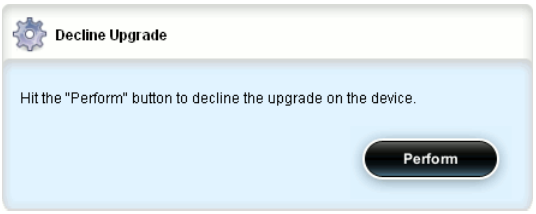

Figure 2.28. Decline Upgrade

To decline an upgrade, click on the *decline-upgrade* menu action and then click the *Perform* button on the Decline Upgrade form.

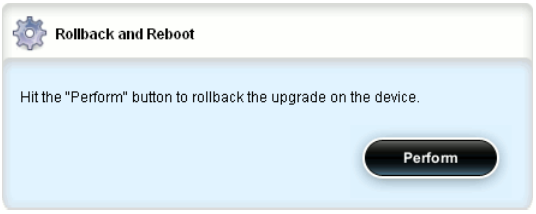

Figure 2.29. Rollback and Reboot

To roll back an upgrade, click on the *rollback-reboot* menu action and then click the *Perform* button on the Rollback and Reboot form.

Rollback and Reboot “rolls back” the system to the previously active software installation, which is stored on the alternate of two filesystem partitions in flash memory. Performing this action will result in rebooting the system using the old software installation along with its configuration.

> *Any configuration changes made since the last software upgrade will not be reflected after rebooting to the "rolled-back" software installation.*

2.6. ROXflash Cross-Partition Imaging Tool - Software Downgrade

ROX™ supports two system partitions. One is always active and the other is inactive. ROXflash allows you to flash any ROX™ software version to the inactive partition.

To obtain a flash image, contact your RuggedCom sales representative. Place the flash image in a location on your network accessible to the ROX™. On the ROXflash form, enter the URL for the flash image and flash it to the inactive partition. The flash image will be active after the next reboot.

2.6.1. Uses

Use ROXflash for downgrading to an earlier version of the ROX software. For example, your organization has certified a specific version of the ROX software, and all ROX™ units must run the certified version. Due to an equipment issue, you need to install a new ROX™ unit that comes with a later version of the software. In this example, use ROXflash to install the earlier version of the software on the new unit.

Use ROXflash only to install earlier versions of the ROX software. Software upgrades to later versions should be performed using the Software Upgrade function.

Table 2.1, “Differences Between ROXflash and Software Upgrade Functions” outlines some of the key differences between the ROXflash and Software Upgrade functions. For more information on the Software Upgrade function, see Section 2.5, “Software Upgrade”.

ROXflash	Software Upgrade
Used primarily for downgrades.	Used only for upgrades; does not support downgrades (except for rollbacks).
Uses a flash image ordered from a RuggedCom Sales Representative.	Uses an archive of ROX™ software packages hosted on an upgrade server. The archive is available on RuggedCom.com for download.
Downgrades to any software version supplied in an image.	Rolls back only to the last version stored on the alternate partition.
Does not transfer system configurations and files to the next software version. ROXflash returns the unit to its factory default settings. Configurations must be reloaded after rebooting.	Transfers configurations and files to the upgraded software version; reverts to the previous configurations in a rolled back version.

Table 2.1. Differences Between ROXflash and Software Upgrade Functions

2.6.2. ROXflash Configuration

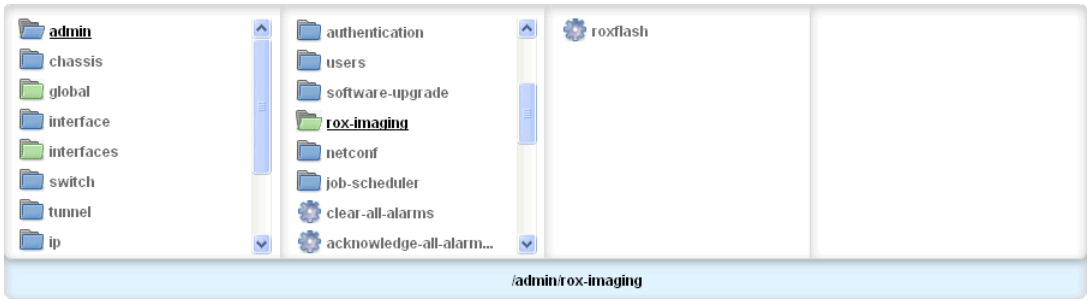

Figure 2.30. ROX-Imaging menu

The ROX-Imaging menu is accessible from the main menu under *admin*. The ROXflash Monitoring form appears on the same screen as this menu.

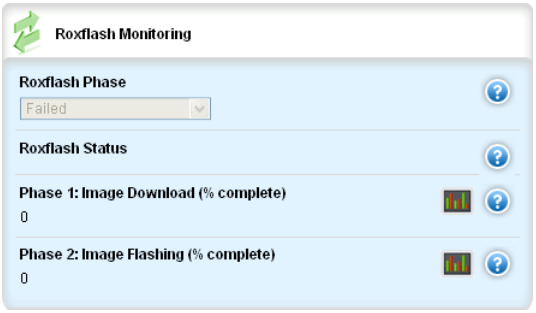

Figure 2.31. ROXflash Monitoring form

This form shows the progress and state of the roxflash operation (during an upgrade or downgrade).

ROXflash Phase

Synopsis: *string - one of the following keywords { Failed, Completed successfully, Unknown state, Imaging partition, Downloading image, Inactive }*

The current phase or state of the ROXflash operation. It is always one of: 'Inactive', 'Downloading image', 'Imaging partition', 'Unknown state', Completed successfully, or 'Failed'. These phrases do not vary, and may be used programatically for ascertaining state.

ROXflash Status

Synopsis: *A string*

Detailed messages about ROXflash progress.

Phase 1: Image Download (% complete)

Synopsis: *integer*

Phase 1 of ROXflash downloads the image from a URL. This reflects percent complete.

Phase 2: Image Flashing (% complete)

Synopsis: *integer*

Phase 2 of ROXflash flashes the image to the alternate partition. This reflects percent complete.

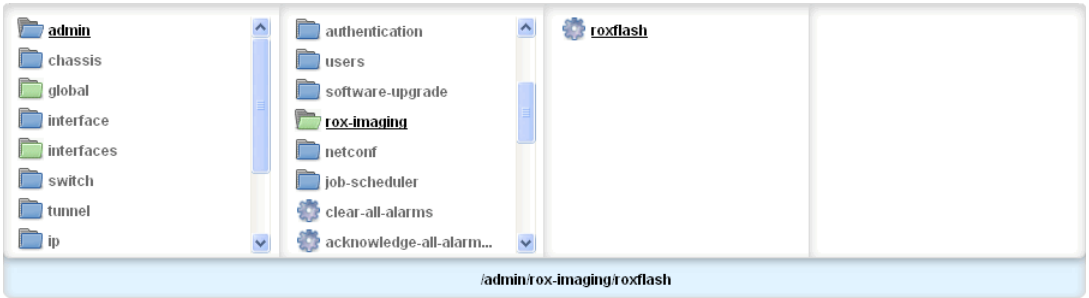

Figure 2.32. ROXFlash menu

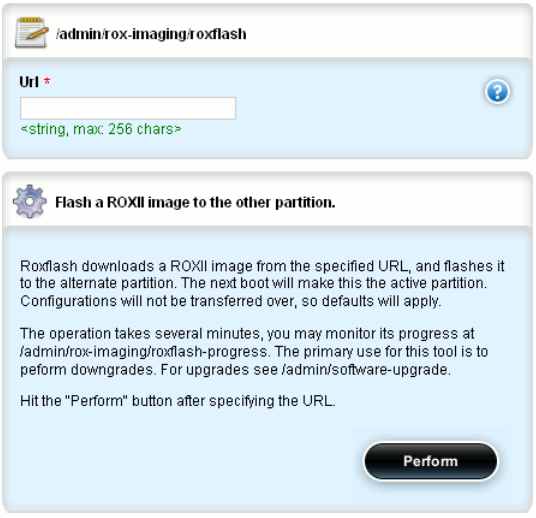
/admin/rox-imaging/roxflash
Url *
<string, max: 256 chars>
Flash a ROXII image to the other partition.
Roxflash downloads a ROXII image from the specified URL, and flashes it to the alternate partition. The next boot will make this the active partition. Configurations will not be transferred over, so defaults will apply.
The operation takes several minutes, you may monitor its progress at /admin/rox-imaging/roxflash-progress. The primary use for this tool is to peform downgrades. For upgrades see /admin/software-upgrade.
Hit the "Perform" button after specifying the URL.
Perform

Figure 2.33. ROXFlash forms

To perform a ROXFlash operation, enter the URL into the ROXflash form and then click the *Perform* button. Next, monitor the progress by returning to the ROXflash Monitoring form.

2.7. Scheduling Jobs

Use job scheduling to execute CLI (command line interface) commands at a specified time and date or in response to configuration changes. The path to the *scheduler* menu is *admin/scheduler*.

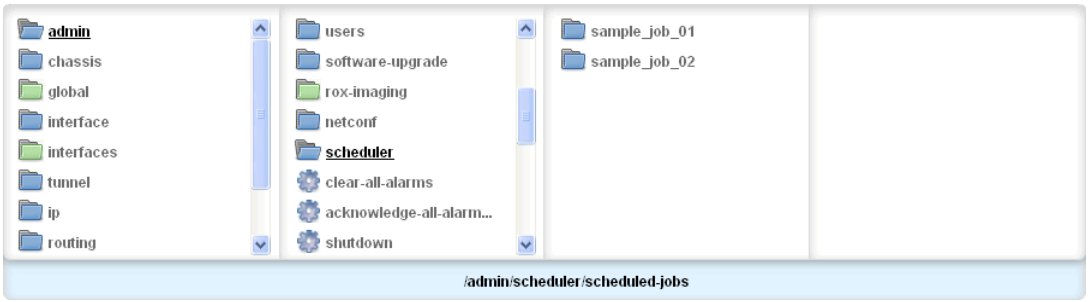

Figure 2.34. Scheduler menu

There are two types of scheduled jobs:

- *periodic* jobs launch at a defined interval. Set the interval in the *Minute*, *Hour*, *Day of Month*, and *Month* parameters. Use the *Day of Week* parameter to launch the job on a specific day of the week, such as every Friday. For information on how periodic scheduled jobs behave when you omit date and time parameters, see Figure 2.36, “Scheduled Jobs Form” and the field descriptions.
- *configchange* jobs launch only when the configuration changes.

The job scheduler *Command* parameter accepts most ROX CLI commands. Do not use commands that require a manual response or confirmation.

The */admin/scheduler/scheduled-jobs* table lists the scheduled jobs and their settings:

/admin/scheduler/scheduled-jobs
Job Name Job Type Minute Hour Day of Month Month Day of Week Command
sample_job_01 periodic 1 not found not found not found monday backupconfig
sample_job_02 periodic 5 5:00 1 1 monday clearmessagelog

Figure 2.35. Scheduled-jobs table

To add a scheduled job:

- Enter edit mode, navigate to *admin/scheduler*, and click **<Add scheduled-jobs>**.
- On the *Key settings* form, enter a name for the job and click **Add**.
- On the *Scheduled Jobs* form, set the job parameters.

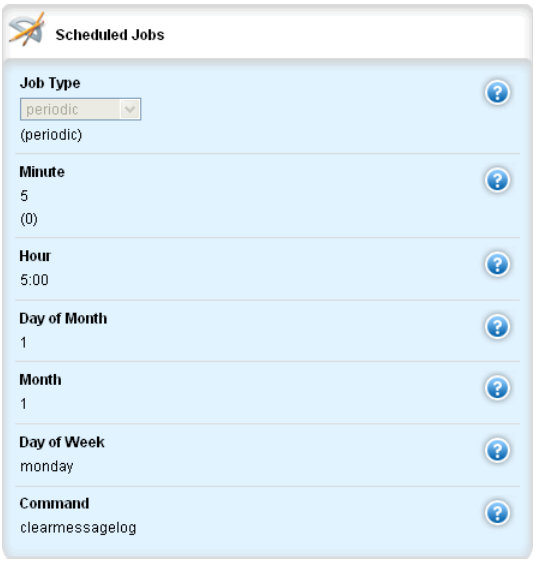

Figure 2.36. Scheduled Jobs form

Job Type

Synopsis: *string - one of the following keywords { periodic, configchange }*

Default: *periodic*

Determines when to launch the scheduled job:

- periodic: the job launches at a set date and time.
- configchange: the job launches when the configuration changes.

Minute

Synopsis: *A string*

Default:

For periodic jobs, sets the minutes portion of the job launch time. Valid values are in the range of 0 to 59. If no value is set, the scheduler uses the default value of 0 and launches the job every hour on the the hour.

- To specify a single value, enter the value in the field. For example, to launch the job 10 minutes past the hour, enter 10
- To specify a list of values, enter the values as a comma-separated list. For example, to launch the job at 14, 30, and 45 minutes past the hour, enter 15,30,45

- To specify a range of values, enter the range as comma-separated values. For example, to launch the job every minute between 30 and 45 minutes past the hour, enter 30-45

Hour

Synopsis: *A string*

For periodic jobs, sets the hour portion of the job launch time, in the 24-hour clock format. Valid values are in the range of 0 to 23. If no value is set, the job launches every hour at the time set in the Minute field.

- To specify a single value, enter the value in the field. For example, to launch the job at 5:00 pm, enter 17
- To specify a list of values, enter the values as a comma-separated list. For example, to launch the job at 9:00 am, 12:00 pm, and 5:00 pm, enter 9,12,17
- To specify a range of values, enter the range as comma-separated values. For example, to launch the job every hour between 9:00 am and 5:00 pm, enter 9-17

Day of Month

Synopsis: *A string*

For periodic jobs, sets the day of the month on which to run the scheduled job. Valid values are in the range of 1 to 31. If no value is set, the job launches every day.

- To specify a single value, enter the value in the field. For example, to launch the job on the tenth day of the month, enter 10
- To specify a list of values, enter the values as a comma-separated list. For example, to launch the job on the first, fifteenth, and thirtieth days of the month, enter 10,15,30
- To specify a range of values, enter the range as comma-separated values. For example, to launch the job on days one through fifteen, enter 1-15

Month

Synopsis: *A string*

For periodic jobs, sets the month in which to run the scheduled job. Valid values are in the rage of 1 to 12. If no value is set, the job launches every day.

- To specify a single value, enter the value in the field. For example, to set the month to February, enter 2
- To specify a list of values, enter the values as a comma-separated list. For example, to set the months to January, June, and December, enter 1,6,12
- To specify a range of values, enter the range as comma-separated values. For example, to set the months to January through June, enter 1-6

Day of Week

Synopsis: *A string*

For periodic jobs, sets the day of the week on which to run the scheduled job. Valid entries are in the range of 0 to 6, where 0 represents Sunday, 1 represents Monday, and so on. If no value is set, the job launches every day.

- To specify a single value, enter the value in the field. For example, to set the day to Monday, enter 1
- To specify a list of values, enter the values as a comma-separated list. For example, to set the days to Friday, Saturday, and Sunday, enter 5,6,0
- To specify a range of values, enter the range as comma-separated values. For example, to set the days to Monday through Friday, enter enter 1-5

Command

Synopsis: *A string*

The CLI commands to execute at the scheduled time. The command or list of commands can be up to 1024 characters in length. For example, this command saves the running configuration to a file named 'myconfig': show running-config | save myconfig

Do not use interactive commands or commands that require a manual response or confirmation.

2.8. The Featurekey

2.8.1. Overview

Some ROX™ software features are only available by purchasing an appropriate feature level. Consult the product datasheet for available feature levels and the specific capabilities they enable.

When specifying a feature level at the time of ordering, the featurekey is entered into the electronic signature on the device . The featurekey is independent of the compact flash card and is retained by the device should the card be replaced.

2.8.2. Upgrading Feature Levels in the field

Feature levels can be purchased and upgraded in the field with a file-based featurekey. To update your featurekey, contact your RuggedCom sales representative. For the RX1000, you need to provide the serial number for the unit you are upgrading. The upgraded featurekey is licensed for the serial number you provide. For instructions on how to view your serial numbers, see Section 2.8.4, “Viewing RuggedCom Serial Numbers”.

To install the featurekey file, use the Install Files form found under that *admin* menu. You can also use the **file scp-featurekey-from-url** command from the ROX™ Command Line Interface. For instructions on how to upload the featurekey file, see Section 2.8.5, “Uploading a Featurekey”.

The upgraded featurekey resides on the device’s compact flash card. ROX™ evaluates both the device featurekey and the file-based featurekey, and then enables the most capable feature level described by the keys.

When using file-based featurekeys, the feature level follows the compact flash card. Moving the compact flash card to another device moves the feature level to the new device. If you want the upgraded feature level to be tied to a specific device, contact your RuggedCom sales representative to arrange for an RMA (Return to Manufacturer Authorization) to have the featurekey programmed into the device.

2.8.3. When a File-based featurekey does not Match the Hardware

In rare circumstances, you may need to remove the compact flash card from one device and transfer it to another device. For example: you may have a backup device to replace a malfunctioning unit, and you choose to use the upgrade featurekey on the malfunctioning unit’s compact flash card to retain your configuration in the backup unit.

The file-based featurekey on the compact flash card is licensed for a particular unit, but can be transferred to another unit to ensure continuity of service. When you transfer the file-based featurekey from its licensed unit to another unit for which it is not licensed, the device behaves in the following manner:

1. The device enables the higher feature level found on the compact flash card.
2. The device raises a non-clearable alarm, indicating a hardware mismatch with the featurekey.
3. The alarm trips the fail-safe relay and turns on the main alarm LED.

To acknowledge the alarm and resolve the issue, follow these steps:

1. Acknowledge the alarm. (For instructions on acknowledging alarms , see Chapter 6, *Alarms*.)
2. Contact a Ruggedcom sales representative and order a featurekey matching the serial numbers of the hardware you are using.

2.8.4. Viewing RuggedCom Serial Numbers

When you order a new featurekey, you need to provide RuggedCom with the chassis serial number. This section describes how to view your device's serial numbers through the **CLI** screen in the ROX™ web interface.

Follow these steps to display the serial numbers for your device:

Procedure 2.1. Viewing RuggedCom Serial Numbers

1. Launch a web browser and navigate to your device's IP address. Log in to ROX™. The ROX web interface appears.
2. Click the **Tools** tab and click the **CLI** link. The CLI screen appears.

Figure 2.37. CLI in the ROX™ Web Interface

3. At the Operational mode command line prompt, type **show chassis** and press **Enter**. Chassis information appears:

```
ruggedcom# show chassis
chassis
 chassis-status
  model RX1000 software license "RX1000 Standard Edition" order code ...
 hardware
  slot-hardware
      ORDER
SLOT  FIELD  DETECTED MODULE                SERIAL NUMBER
-------------------------------------------------------------
pm1   48     48VDC (36-59VDC) Power Supply
lm1   XX     none                          none
lm2   M1_    Old V90 Modem
lm3   TX01   2x 10/100Tx RJ45
lm4   TX01   2x 10/100Tx RJ45
lm5   DS3    1x T3/E3
lm6   TC2    2x Chan T1/E1
pm2   XX     none                          none
main  CM01   RX1000 Main Board             RX1K-12-11-0015
```

In the **`slot-hardware`** table, make note of the **`main`** slot serial number (highlighted in **bold text** in the example above).

4. When ordering a new featurekey, provide the **`main`** slot serial number to RuggedCom.

2.8.5. Uploading a Featurekey

After receiving your featurekey file from RuggedCom, save the file to a computer that is accessible to your device through your network.

2.8.5.1. Uploading a Featurekey Using the Web User Interface

Install Featurekey files using the Install Files forms found under the *admin* menu.

To install a featurekey file, navigate to *admin/install-files*. The Install Files form appears. In the in the *File type* field, select **featurekey**. In the *URL* field, enter the URL to the file. On the *Install Files to Devices* form, click the *Perform* button.

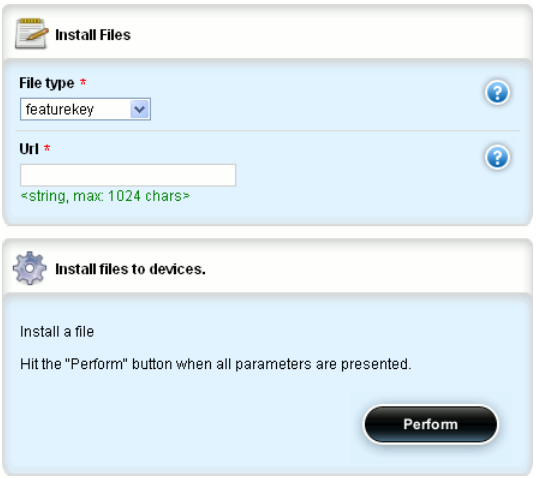

Figure 2.38. Install Files forms

For more information on installing files, see Section 2.9.1, “Installing Files”.

2.8.5.2. Uploading a Featurekey Using the Command Line Interface

To upload the file to your device, you will need to know the following information:

- the featurekey filename.
- a user name and password to log in to the computer where you saved the featurekey file.
- the hostname or IP address of the computer where you saved the featurekey file.

Follow these steps to upload a featurekey file to your device:

Procedure 2.2. Uploading a Featurekey File

1. Launch a web browser and navigate to your device’s IP address. Log in to ROX™. The ROX™ web interface appears.
2. Click the **Tools** tab and click the **CLI** link. The CLI screen appears.
3. In Operational mode, at the command line prompt, type the following command:

   file scp-featurekey-from-url {username}@{host}:{path}{filename} {filename}

   Where:

   - **{username}** is the name of a user who can log in to the computer where you saved the featurekey file.
   - **{host}** is hostname or IP address of the computer where you saved the featurekey file.
   - **{path}** is the directory path to the featurekey file on the computer.

- **{filename}** is the name of the featurekey file.

For example:

file scp-featurekey-from-url wsmith@10.200.20.39:/files/keys/1_cmRX1K-12-11-0015.key 1_cmRX1K-12-11-0015.key

4. Type the command with your parameters and press **Enter**. When prompted, type the user's password and press **Enter**. The system uploads the featurekey file:

```
ruggedcom# file scp-featurekey-from-url wsmith@10.200.20.39:/files/keys/
1_cmRX1K-12-11-0015.key 1_cmRX1K-12-11-0015.key
wsmith@10.200.20.39's password:
1_cmRX1K-12-11-0015.key               100%  192     0.2KB/s   00:00
ruggedcom#
```

5. To view the contents of the featurekey file, type the following command:

file show-featurekey {filename}

Where:

- **{filename}** is the name of the featurekey file.

For example:

file show-featurekey 1_cmRX1K-12-11-0015.key

6. Type the command with your featurekey filename and press **Enter**. The system displays the contents of the featurekey file:

```
ruggedcom# file show-featurekey 1_cmRX1K-12-11-0015.key
GPG_FEATUREKEY_LEVEL=1
GPG_FEATUREKEY_CM_SERIALNUMBER=RX1K-12-11-0015
GPG_FEATUREKEY_SIGNATURE=iEYEABECAAYFAk091pAACgkQP2pya+G5kdZeKACeKdHUB2G1T73Dymq8IjSdYDKAiskAn
3abBpCEhfLXxY2ZlVbvGNwDZow2
ruggedcom#
```

7. On the CLI screen, click **Stop** to close the CLI session.

2.8.6. Backing Up a Featurekey Using the Web User Interface

Featurekey files can be backed up using the following forms. These forms are accessible from the *admin* menu.

To back up a featurekey file, navigate to *admin/backup-files*. The Backup Files form appears. In the *File type* field, select **featurekey**. Enter additional parameters on the form. On the *Backup Files From Devices* form, click the *Perform* button.

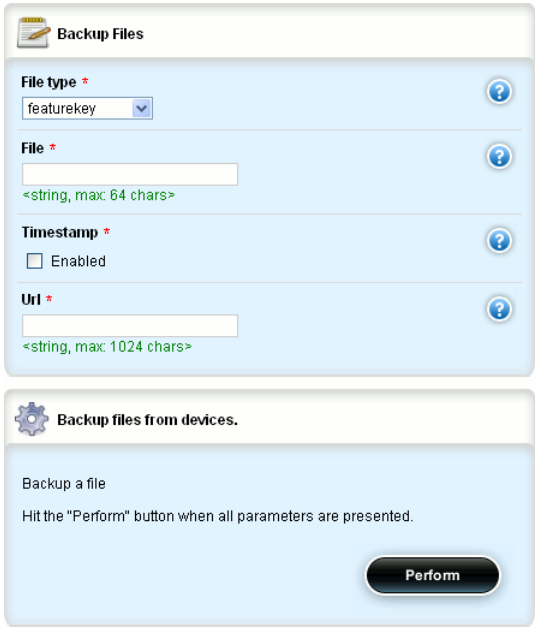

Figure 2.39. Backup Files forms

For more information on backing up files, see Section 2.9.2, “Backing Up Files”.

2.9. Installing and Backing Up Files

You can install and back up files using the following forms found under the *admin* menu.

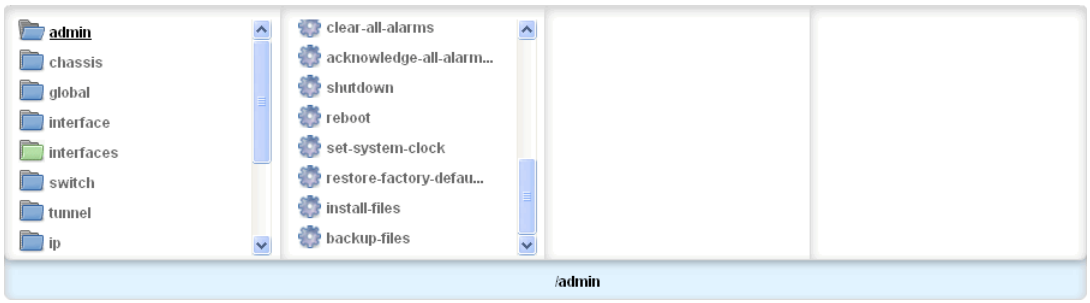

Figure 2.40. Administration menu

2.9.1. Installing Files

To install a file, click *install-files*. The Install Files forms appear.

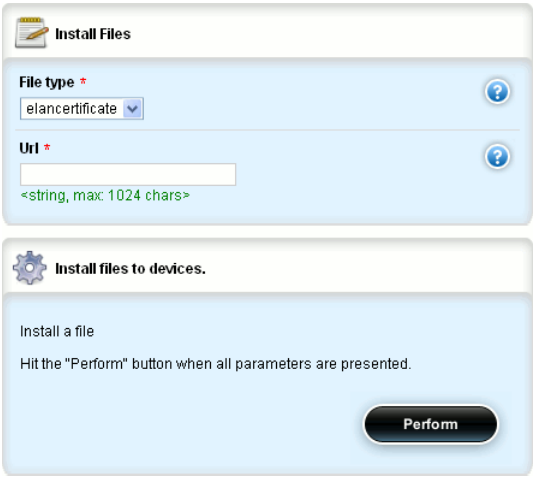

Figure 2.41. Install Files forms

On the *Install Files* form, select the file type and enter a URL. On the Install Files To Devices form, click the *Perform* button.

2.9.2. Backing Up Files

To back up a file, click on *backup-files*. The Backup Files forms appear.

Figure 2.42. Backup Files forms

On the *Backup Files* form, select the file type and enter the required parameters. On the Backup Files From Devices form, click the *Perform* button.

2.10. Deleting Log Files

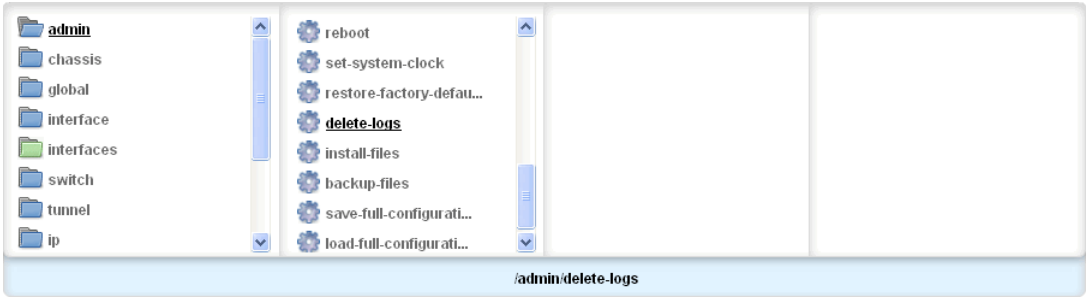

Figure 2.43. Delete-logs menu

To delete log files, click the **Perform** button on the Delete Log Files form. This form is accessible at *admin/delete-logs*.

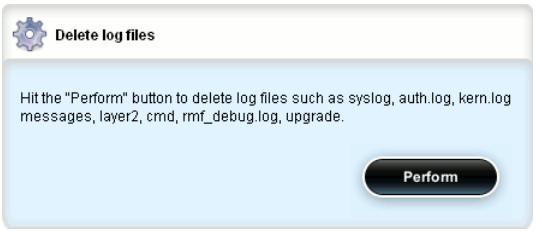
Delete log files
Hit the "Perform" button to delete log files such as syslog, auth.log, kern.log messages, layer2, cmd, rmf_debug.log, upgrade.
Perform

Figure 2.44. Delete Log Files form

2.11. Saving Full Configurations

Save full configurations to a file using the forms below. These forms are accessible at *admin/save-full-configuration*.

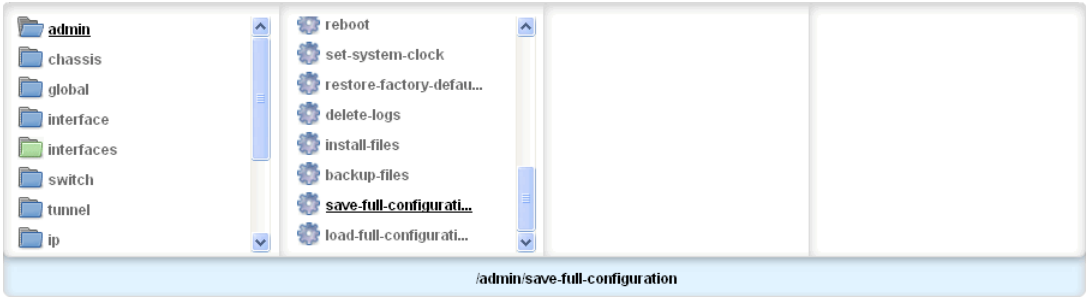

Figure 2.45. Save-full-configuration menu

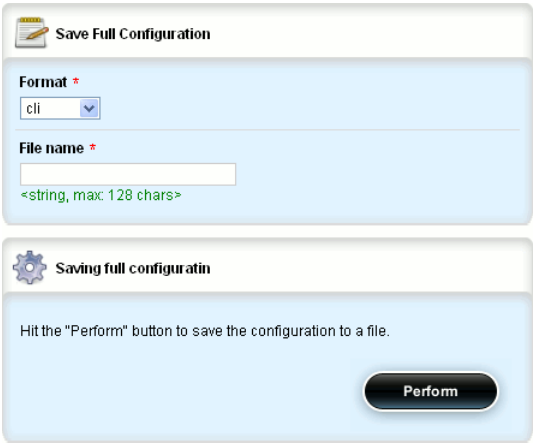

Figure 2.46. Save Full Configuration forms

To save full configurations to a file, select the format and enter the parameters in the Save Full Configuration form, then click the **Perform** button in the Saving Full Configuration form.

2.12. Loading Full Configurations

Load full configurations to a file using the forms below. These forms are accessible at *admin/load-full-configuration*.

Figure 2.47. Load-full-configuration menu

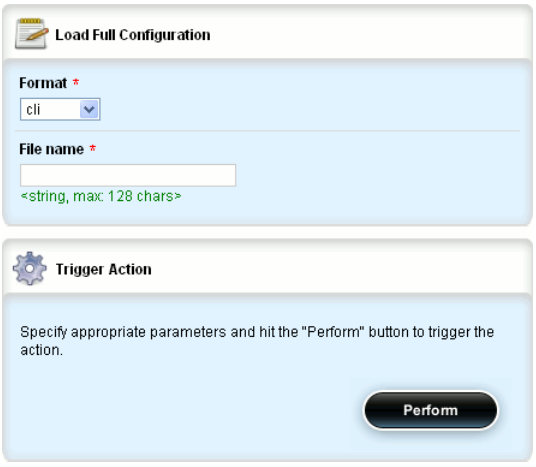

Figure 2.48. Load Full Configuration forms

To load full configurations to a file, select the format and enter the parameters in the Load Full Configuration form, then click the **Perform** button in the Trigger Action form.

3. Time Synchronization

ROX™ offers the following timekeeping and time synchronization features:

- Local hardware timekeeping and time zone management
- NTP time synchronization

3.1. NTP Fundamentals

NTP (Network Time Protocol) is an Internet protocol used to synchronize the clocks of computers to some time reference. Variants of NTP such as SNTP (Simple NTP, a reduced functionality NTP) and XNTP (Experimental NTP) exist. NTP itself is available in versions 3 and 4 (RuggedRouter® includes version 4).

NTP is a fault-tolerant protocol that allows an NTP daemon program to automatically select the best of several available time sources, or reference clocks, to synchronize to. Multiple candidates can be combined to minimize the accumulated error. Temporarily or permanently wrong time sources are detected and avoided.

The NTP daemon achieves synchronization by making small and frequent changes to the router hardware clock.

The NTP daemon operates in a *client-server mode*, both synchronizing from *servers* and providing synchronization to *peers*.

If NTP has a number of servers to choose from, it will synchronize with the lowest stratum server. The stratum is a measure of the number of servers to the (most highly accurate) reference clock. A reference clock itself appears at stratum 0. A server synchronized to a stratum n server will be running at stratum n + 1.

You will generally configure lower stratum NTP hosts as servers and other NTP hosts at the same stratum as peers. If all your configured servers fail, a configured peer will help in providing the NTP time. It is generally a good idea to configure one at least one server and peer.

The NTP daemon will know about the NTP servers and peers to use in three ways.

- It can be configured manually with a list of servers to poll,
- It can be configured manually with a list of peers to send to,
- It can look at advertisements issued by other servers on multicast or broadcast addresses.

Note that if multicasting or broadcasting is used, it is strongly recommended to enable authentication unless you trust all hosts on the network.

NTP uses UDP/IP packets for data transfer because of the fast connection setup and response times UDP offers. The NTP protocol uses port UDP port 123. Note that if your router employs a firewall and acts as a client it must open UDP port 123. Additionally, if the router acts as a server the firewall must allow connection requests on port 123 as well.

3.1.1. The NTP Sanity Limit

The NTP daemon corrects the system time through two means, “stepping” and “slewing”. If the difference between the local clock and the reference clock chosen by NTP (the “offset”) is more than 128ms for a period of more than 900 seconds, NTP will “step” or instantaneously correct the time. If the time difference is less than 128ms, NTP will “slew” the time by no more than 500 microseconds every second towards the correct time, in such a way that to an application on the system, the time never appears to be flowing backwards.

After booting, NTP uses slewing to achieve synchronization by making small and frequent changes to the router hardware clock. If the reference server's clock differs from the local clock by more than 1000 seconds, the NTP daemon decides that a major problem has occurred and terminates.

3.2. Configuring Time Synchronization

To configure time synchronization, configure the following items:

- set the system time and date. See Section 3.2.1, "Configuring the System Time and Date".
- set the system timezone. See Section 3.2.2, "Configuring the System Time Zone".
- set the local time settings. See Section 3.2.3, "Configuring the Local Time Settings".
- add remote NTP servers. You can add remote NTP servers with or without authentication. See Section 3.2.4, "Configuring NTP Servers".
- set the NTP server restrictions. See Section 3.2.6, "Configuring NTP Server Restrictions".
- configure an NTP server using Multicast or Broadcast. See Section 3.2.7, "Configuring an NTP Server using Multicast or Broadcast".
- configure an NTP client using Multicast. See Section 3.2.8, "Configuring an NTP Client using Multicast".
- configure an NTP client using Broadcast. See Section 3.2.9, "Configuring an NTP Client using Broadcast".

After configuring NTP, you can check the status of the NTP service. See Section 3.2.10, "Checking NTP Status".

3.2.1. Configuring the System Time and Date

To set the system time and date:

- Navigate to *admin/set-system-clock*.
- On the *Set New Time and Date* form, enter the date in the format YYYY-MM-DD HH:MM:SS.

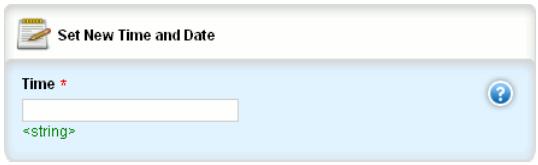

Figure 3.1. Set new Time and Date form

- On the *Set clock on target device* form, click **Perform**.

3.2.2. Configuring the System Time Zone

To set the system time zone:

- In edit mode, navigate to *admin*.
- On the *Timezone* form, select a timezone from the list.

  The Etc/GMT timezones conform to the POSIX style and have their signs reversed from common usage. In POSIX style, zones west of GMT have a positive sign; zones east of GMT have a negative sign.

Figure 3.2. Timezone form

- Commit the changes.

3.2.3. Configuring the Local Time Settings

On the *Local Time Settings* form, you enable the local clock and set the NTP stratum level.

The path to the *Local Time Settings* form is */services/time/ntp*.

To set the local time settings:

- In edit mode, navigate to */services/time/ntp*.
- On the *Local Time Settings* form, set the local time parameters.
- Commit the changes.

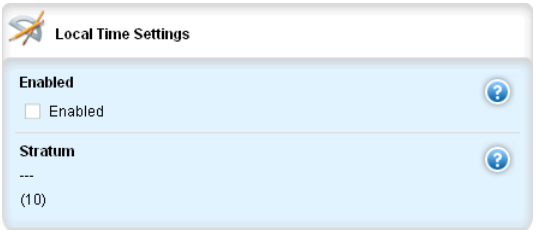

Figure 3.3. Local Time Settings form

Enable

Enables the local clock

Stratum

Synopsis: *unsigned byte integer*

Default: *10*

The stratum number of the local clock

3.2.4. Configuring NTP Servers

ROX™ can periodically refer to an NTP server to correct any accumulated drift in the onboard clock. ROX™ can also serve time via SNTP to hosts that request it.

You can add NTP servers with or without authentication keys. To associate an authentication key with an NTP server, you must first define the server key. For instructions on how to create server keys, see Section 3.2.5, “Adding Server Keys”.

To view the list of configured NTP servers, navigate to */services/time/ntp/server*.

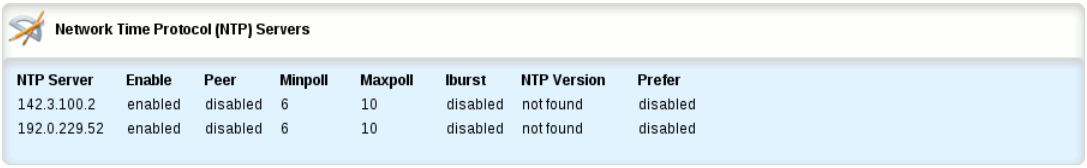
Network Time Protocol (NTP) Servers
NTP Server Enable Peer Minpoll Maxpoll Iburst NTP Version Prefer
142.3.100.2 enabled disabled 6 10 disabled not found disabled
192.0.229.52 enabled disabled 6 10 disabled not found disabled

Figure 3.4. Network Time Protocol (NTP) Servers

To add an NTP server:

- In edit mode, navigate to */services/time/ntp/server* and click **<Add server>**.
- On the *Key settings* form, enter the IP address or hostname for the server and click **Add**.
- On the *Network Time Protocol (NTP) Servers* form, set the server parameters.
- Commit the changes.

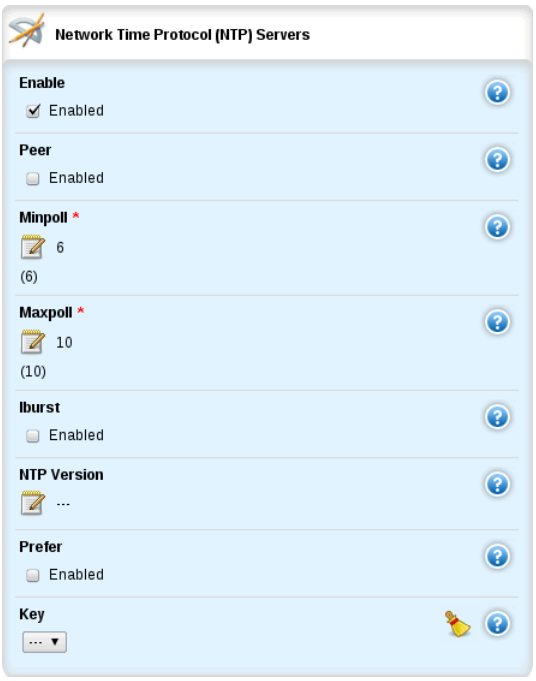

Figure 3.5. Network Time Protocol (NTP) Servers form

Enable

Turns on the NTP interface to this server.

Peer

Allows you to enter and edit peers. Peers are NTP servers of the same stratum as the router, and are useful when contact is lost with the hosts in the NTP servers menu.

Minpoll

Synopsis: *unsigned byte integer*

Default: *6*

Minimum poll interval for NTP messages, in seconds as a power of two.

Maxpoll

Synopsis: *unsigned byte integer*

Default: *10*

Maximum poll interval for NTP messages, in seconds as a power of two.

Iburst

When the server is unreachable and at each poll interval, send a burst of eight packets instead of the usual one.

NTP Version

Synopsis: *integer*

The version of the NTP protocol used to communicate with this host. Change this only if it is known that the host requires a version other than 4.

Prefer

Marks this server as preferred.

Key

Synopsis: *unsigned short integer*

An authentication key associated with this host.

3.2.5. Adding Server Keys

Use server keys to use authentication for NTP communications. NTP authentication authenticates the time source to help prevent tampering with NTP timestamps. When using authentication, both the local and remote servers must share the same key and key identifier. Packets sent to and received from the server/peer include authentication fields encrypted using the key.

Keys defined here are associated with NTP servers on the *Network Time Protocol (NTP) Servers* and *NTP Broadcast/Multicast Servers* forms.

To add a server key:

- In edit mode, navigate to */services/time/ntp/key* and click **<Add key>**.
- On the *Key settings* form, enter an identifier for the key and click **Add**.
- On the *Server Keys* form, set the key parameters.
- Commit the changes.

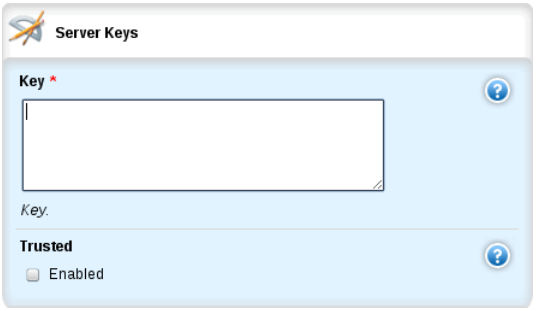

Figure 3.6. Server Keys form

Key

Synopsis: *"AES CFB128"-encrypted string*

Key.

Trusted

Mark this key is trusted for the purposes of authenticating peers with symmetric key cryptography. The authentication procedures require that both the local and remote servers share the same key and key identifier.

3.2.6. Configuring NTP Server Restrictions

Use server restrictions to control and restrict access to the NTP server.

To set NTP server restrictions:

- In edit mode, navigate to */services/time/ntp/restrict* and click **<Add restrict>**.
- On the *Key settings* form, set the following parameters and click **Add**.

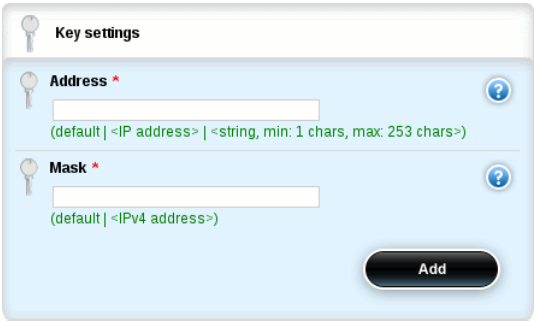

Figure 3.7. Server Restrictions Key settings form

Address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Synopsis: *Domain name (RFC 1034)*

Synopsis: *string - the keyword { default }*

Address to match. The address can be host or network IP address or a valid host DNS name.

Mask

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *string - the keyword { default }*

Mask used to address match. Mask 255.255.255.255 means address is treated as the address of an individual host.

- On the *Server Restrictions* form, set the restriction parameters.
- Commit the changes.

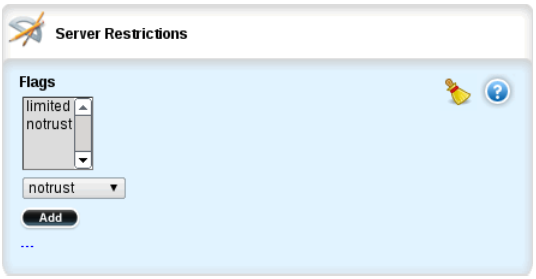

Figure 3.8. Server Restrictions form

Flags

Synopsis: *string - one of the following keywords { version, ntpport, notrust, notrap, noserve, noquery, nopeer, nomodify, lowpriotrap, limited, kod, ignore }*

Synopsis: *"flags" occurs in an array.*

Flags restrict access to NTP services. An entry with no flags allows free access to the NTP server.

- version: denies packets that do not match the current NTP version.
- ntpport: matches only if the source port in the packet is the standard NTP UDP port (123).
- notrust: denies service unless the packet is cryptographically authenticated.
- notrap: declines to to provide mode 6 control message trap service to matching hosts.
- noserve: denies all packets except ntpq(8) and ntpdc(8) queries.
- noquery: denies ntpq(8) and ntpdc(8) queries.

- nopeer: denies packets which result in mobilizing a new association.
- nomodify: denies ntpq(8) and ntpdc(8) queries attempting to modify the state of the server; queries returning information are permitted.
- lowpriotrap: declares traps set by matching hosts to be low priority.
- limited: denies service if the packet spacing violates the lower limits specified in the NTP discard setting.
- kod: sends a kiss-o-death (KoD) packet when an access violation occurs.
- ignore: denies all packets.

3.2.7. Configuring an NTP Server using Multicast or Broadcast

The NTP broadcast/multicast address must be the same as the client address. It is recommended that NTP authentication be used and that a server key be set with the broadcast/multicast setting. For instructions on how to create server keys, see Section 3.2.5, “Adding Server Keys”.

To set a multicast/broadcast address for an NTP server:

- In edit mode, navigate to */services/time/ntp/broadcast* and click **<Add broadcast>**.
- On the *Key settings* form, enter the broadcast/multicast IP address and click **Add**.
- On the *NTP Broadcast/Multicast Servers* form, set the broadcast/multicast parameters.
- Commit the changes.

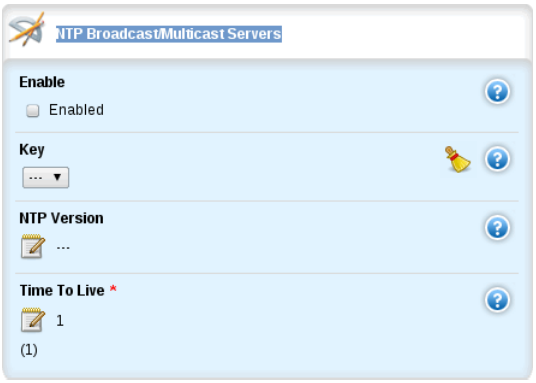

Figure 3.9. NTP Broadcast/Multicast Servers form

Enable

Enables sending broadcast or multicast NTP messages to this address.

Key

Synopsis: *unsigned short integer*

Authentication key.

NTP Version

Synopsis: *integer*

The version of the NTP protocol used to communicate with this host. Change this only if it is known that the host requires a version other than 4.

Time To Live

Synopsis: *unsigned byte integer*

Default: *1*

Time to live.

3.2.8. Configuring an NTP Client using Multicast

Configuring a multicast address for an NTP client enables the client to listen for and receive NTP messages on the multicast address. It is recommended that NTP authentication be used and that a server key be set with the multicast setting. For instructions on how to create server keys, see Section 3.2.5, “Adding Server Keys”.

To set a multicast address for an NTP client:

- In edit mode, navigate to */services/time/ntp*.
- On the *NTP Multicast Clients* form, set the multicast parameters.
- Commit the changes.

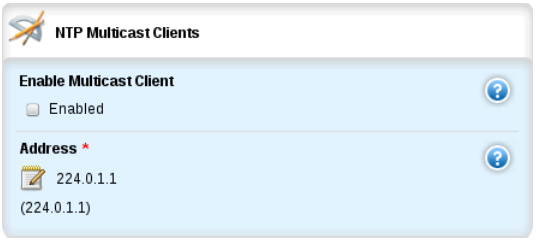

Figure 3.10. NTP Multicast Clients form

Enable Multicast Client

Enables the multicast message mode

Address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Synopsis: *Domain name (RFC 1034)*

Default: *224.0.1.1*

The multicast address on which the NTP client listens for NTP messages.

3.2.9. Configuring an NTP Client using Broadcast

Configuring a broadcast address for an NTP client enables the client to listen for and receive NTP messages on the broadcast address, and enables the NTP server to send NTP messages on the broadcast/multicast address. It is recommended that NTP authentication be used and that a server key be set with the broadcast setting. For instructions on how to create server keys, see Section 3.2.5, “Adding Server Keys”.

To set a broadcast address for an NTP client:

- In edit mode, navigate to */services/time/ntp*.
- On the *Network Time Protocol (NTP)* form, set the broadcast parameters.
- Commit the changes.

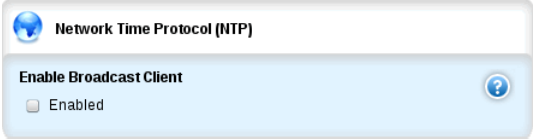

Figure 3.11. Network Time Protocol (NTP) form

Enable Broadcast Client

The broadcast address on which the NTP client listens for NTP messages.

3.2.10. Checking NTP Status

To view the NTP service status:

- In normal or edit mode, navigate to */services/time/ntp/ntp-status* and click <ntp-status>.
- On the *Trigger Action* form, click **Perform**.
- Review the NTP service status in the *NTP Service Status* form.

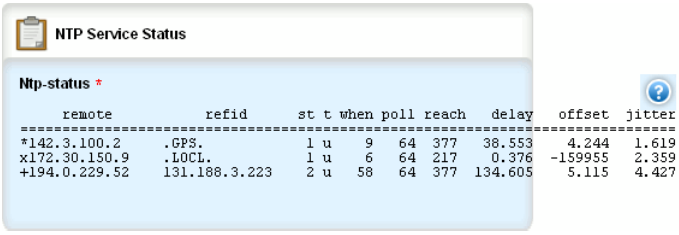

Figure 3.12. NTP Service Status form

For more information on viewing NTP status information, refer to http://support.ntp.org/bin/view/Support/TroubleshootingNTP

4. Basic Network Configuration

This chapter discusses the following:

- IP Interfaces
- Configuring IPv4 and IPv6 Addresses
- Simple Network Setups with IPv4 and IPv6 Addresses

4.1. IP Interfaces

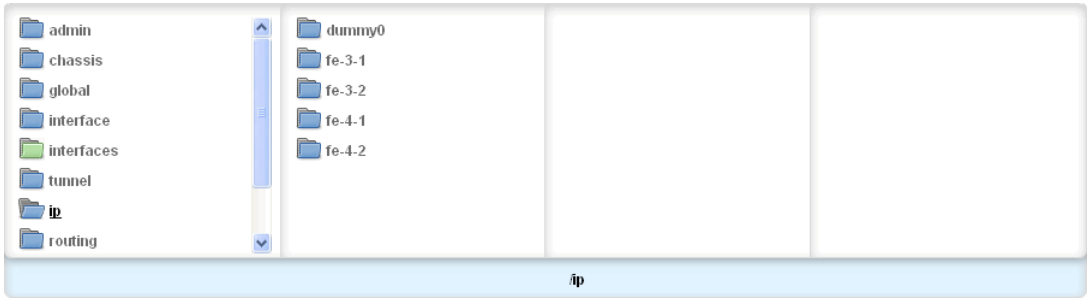

Figure 4.1. IP menu

The IP menu is accessible from the main menu under *ip*.

4.1.1. Configuring an IP Address

The RX1000 has a default IP address configured for each interface. Addresses are configured under the *ipv4* submenu. The first IP address ends with 0.1, with the addresses incrementing for each interface. For example: 192.168.0.1/24, 192.168.1.1/24, 192.168.2.1/24, 192.168.3.1/24.

To configure a different IP address, see Procedure 4.1, “Configuring an IP Address”.

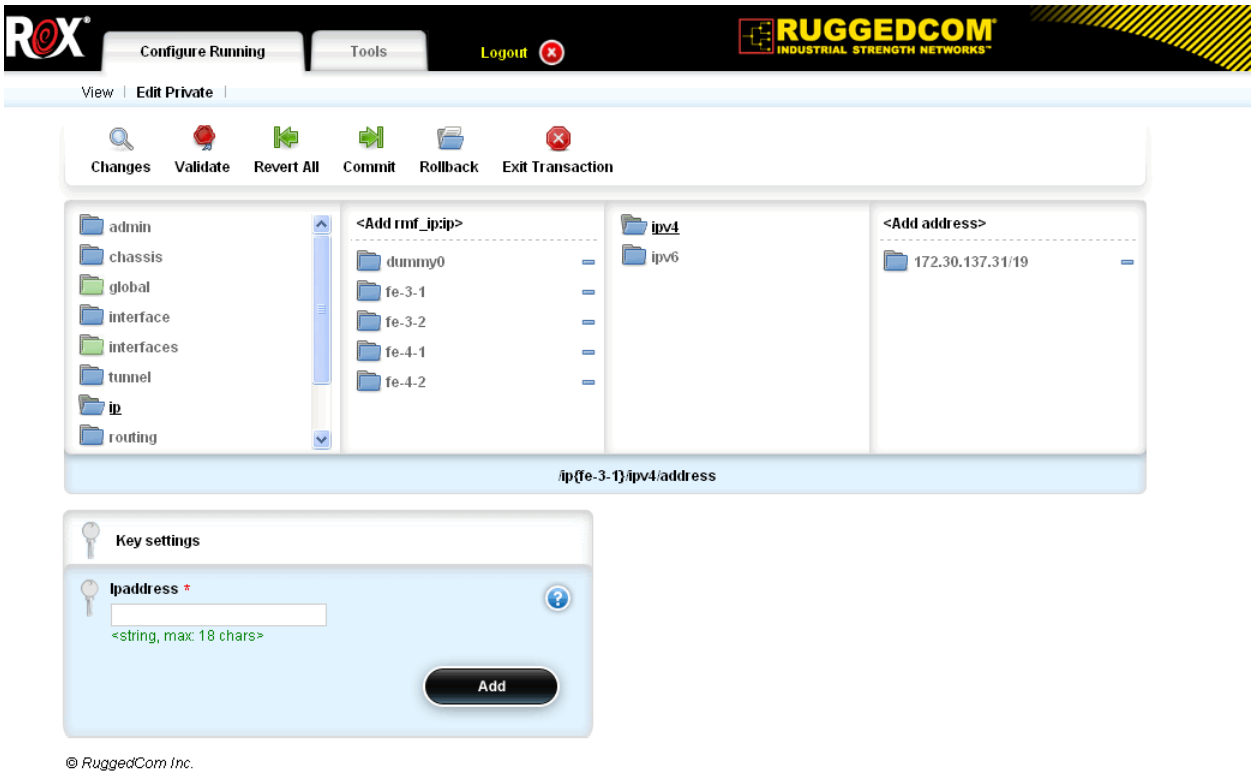

Figure 4.2. Configuring an IP Address

Procedure 4.1. Configuring an IP Address

1. Enter Edit Private mode.
2. Navigate to *ip/interface/ipv4*.
3. To delete an existing IP address, click the *delete* icon.
4. Click *Add address*. The Key settings form appears.
5. In the *IPaddress* field, type the new IP address.
6. Click **Commit**.
7. Click **Exit Transaction**.

4.1.2. Simple Network Setup with the Default IPv4 Addresses

This section describes how to set up a simple network using the factory default IPv4 address.

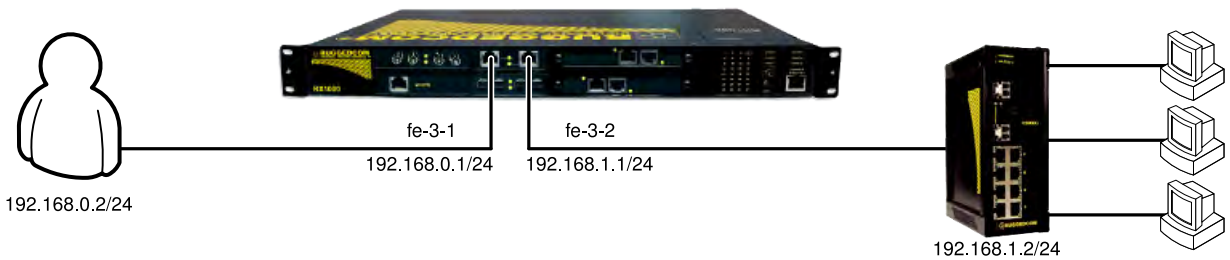

Figure 4.3. Basic Network Setup Using the Default IPv4 Addresses

Procedure 4.2. Basic Network Setup Using the Default IPv4 Addresses

1. Connect a user PC to the Fast Ethernet port (fe-3-1) of the RX1000 and configure the PC to be on the same subnet as the port.
2. Configure the PC to use the IP address, fe-3-1, as the default gateway
3. Connect port fe-3-2 of the RX1000 to a switch typically connecting a LAN
4. The PCs connected to the switch should be on the same subnet as the switch.
5. Configure the switch and the PCs behind the switch to use fe-3-2's IP address (192.168.1.1) as the default gateway
6. From the user PC, ping the IP addresses of the PCs behind the switch. Verify the ping is successful.

To configure a WAN port and assign an IP address, see Chapter 20, *WAN*.

To configure Dynamic Routing on the unit, see Chapter 22, *Dynamic Routing*.

To configure Static Routes and Default Gateways, see Chapter 23, *Static Routing*.

For information related to the Firewall and IP NAT that might be necessary before connecting the unit to the INTERNET, see Chapter 26, *Firewall*.

For information on Dynamic IP address assignment and ProxyARP on non-switched ports, see Section 5.3.1, “Configuring IP Address Source and ProxyARP for Non-switched Interfaces”.

4.1.3. Configuring an IPv6 Address

IPv6 link local addresses starting with the prefix FE80 are assigned to all routable Ethernet interfaces in the RX1000. The Link Local addresses are hidden in the Web UI but they are visible from the CLI (Command Line Interface) using the **show interfaces ip** command.

To advertise IPv6 link layer addresses to their neighbors on the same link, IPv6 Router Advertisement in IPv6 Neighbor Discovery must be enabled. For more information on IPv6 fundamentals and Neighbor Discovery, see Section 5.1, “IPv6 Fundamentals” and Section 5.2, “IPv6 Neighbor Discovery”.

Procedure 4.3. Configuring an IPv6 Address

1. Enter Edit Private mode.
2. From the WEB UI Navigate to ip/interface/ipv6.
3. Click Add address. The Key settings form appears.
4. In the IPaddress field, type an IPv6 address with a network prefix
5. Click Commit.
6. Click Exit Transaction.
7. To delete an existing IPv6 address, click the delete icon under ip/interface/ipv6.
8. Refer to steps 3 to 7 to configure a new IPv6 address

4.1.4. Simple Network Setup with IPv6 Addresses

This section describes how to configure a simple network using the factory default IPv6 address.

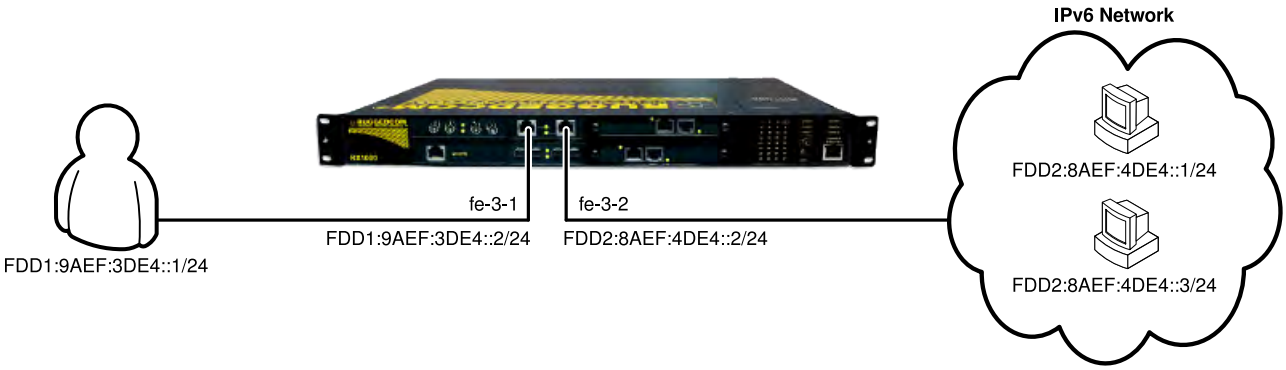

Figure 4.4. Simple IPv6 Network Setup

Procedure 4.4. Simple IPv6 Network Setup

1. Connect a user PC to Fast Ethernet port (fe-3-1) of the RX1000 and configure the PC to be on the same subnet as the port.
2. Configure the S.PC with IPv6 address FDD1:9AEF:3DE4::1/24 and Default Gateway as FDD1:9AEF:3DE4::2.
3. Configure the fe-3-1 and fe-3-2 interfaces of the RX1000 with the IPv6 addresses shown in Figure 4.4, “Simple IPv6 Network Setup”.
4. Connect one of the switched ports from any available LMs to an IPv6 capable network.
5. Configure the D.PCs on the IPv6 network to be on the same IP subnet as fe-3-2 and configure the Default Gateway address as FDD2:8AEF:4DE4::2/48.
6. Enable IPv6 Neighbor Discovery under ip/{interface}/ipv6/nd. For more information on IPv6 neighbor discovery, see Section 5.2, “IPv6 Neighbor Discovery”.
7. Confirm that you can reach the D.PCs from the S.PC.

4.1.5. Routable Interfaces

Figure 4.5. Routable Interfaces table

The Routable Interfaces table is accessible from the *ip* menu.

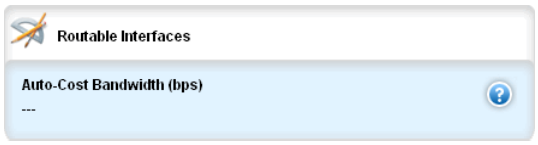

Figure 4.6. Routable Interfaces form

The path to the Routable Interfaces form is *ip/{interface}*.

Interface Name

Synopsis: *A string*

The name for this routable logical interface

Auto-Cost Bandwidth (kbps)

Synopsis: *unsigned long integer*

This value is used in auto-cost calculations for this routable logical interface in kbps

Figure 4.7. Addresses table

The path to the Addresses table is *ip/{interface}/ipv4*. The Addresses table provides a summary of which IP addresses are configured.

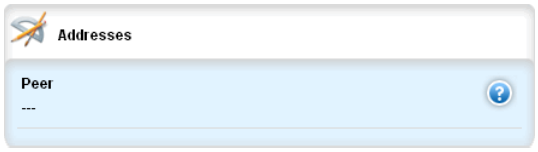

Figure 4.8. Addresses form

The path to the Addresses form is *ip/{interface}/ipv4/{address}*.

ipaddress

Synopsis: *IPv4 address and prefix in CIDR notation*

The IPv4/Prefix (xxx.xxx.xxx.xxx/xx).

peer

Synopsis: *IPv4 address in dotted-decimal notation*

The peer IPv4 Address (xxx.xxx.xxx.xxx, PPP link only).

5. IP Network Interfaces

This chapter familiarizes the user with:

- IPv6 Fundamentals and IPv6 Neighbor Discovery
- Adding VLAN Interfaces to Switched Ports
- Configuring IP Address Source and ProxyARP for Switched and Non-switched Interfaces

5.1. IPv6 Fundamentals

Version 6 of the Internet Protocol (IPv6, RFC 2460) has been designated to replace IPv4 throughout the Internet. Some important changes that IPv6 introduces relative to IPv4 fall into the following categories:

5.1.1. Addressing

IPv6 addresses are four times the length of IPv4 addresses, at 128 bits, to be used as 64 bits of network and 64 bits of host address. The larger address space allows much greater flexibility in hierarchical network definition and routing.

The IPv6 packet header has been simplified relative to IPv4 in order to simplify and therefore speed the processing of packets by routing nodes. It also features more efficiently encoded options and greater flexibility in creating extensions.

5.1.2. Security

Security has been designed into IPv6, rather than being treated as a component that must be added to existing IPv4 network stacks.

5.1.3. IPv6 Address Scopes

There are three scopes of IPv6 addresses named Link Local, Unique Local and Global. A Link Local address is automatically assigned to any IPv6 capable interface. This address is mandatory for the devices on the same link to communicate with each other.

The link local address begins with “FE80” in the first 10 bits of an IPv6 address and the address is not routable. The scope for Unique Local address is within enterprise networks. It identifies the boundary of private networks within an organization. Example of a link local address: **FE80:0000:0000:0000:020A:DCFF:FE01:0CCD**

Unique Local addresses are similar to private IPv4 addresses and they are not routable on the Internet. A Unique Local address consists of the first 7 bits as the site address starts with “FD”, the next 1 bit set to 1 meaning locally assigned, next 40 bits as the Global ID to identify a company, next 16 bits as the Subnet ID to identify the subnets within a site and it is usually defined based on hierarchical plan, and finally the last 64 bits for the Interface ID. Example of a unique local address: **FD00:ABAB:CDCD:EFEF:020A:DCFF:FE01:0CCD**

The Global IPv6 addresses are routable and they are interned to be used on the Internet. In order to allow address aggregation the global addresses are structured in hierarchical order. A global address is identified by the first 48 bits specified by the service provider as the global routing prefix in which the first 3 bits of the address start with 001 (2000::/3), the next 16 bits after the global routing prefix are used to define subnets and the last 64 bits are used for Interface ID to define a host. Example of a unique local address: **2001:0CCD:3456:789A:8A9C:BCAB:023A:1234**

5.1.4. IPv6 Multicast Addresses

In IPv6 multicast addresses are widely used. The use of broadcast address is removed in IPv6, instead IPV6 multicast addresses are used for neighbor discovery and route advertisement. An IPv6 multicast address starts with first 8 bits all set to 1 (FF), next 4 bits to define the Lifetime (0 - Permanent, 1 -

Temporary), then the following 4 bits to define the scope (1 - Node, 2 - Link, 5 - Site, 8 – Organization and E – Global) and the last 112 bits identify a multicast Group ID. Some well-known multicast addresses are mentioned below:

IPv6 M.Cast Address	Scope	Description
FF02::1	Link-Local	All Nodes on a Link
FF02::2	Link-Local	All Routers on a Link
FF01::1	Node-Local	Same Node
FF01::2	Node-Local	Same Router
FF05::2	Site-Local	All Routers on a Site
FF02::1:FFxx:xxxx	Link-Local	Solicited Node Address

Table 5.1. Multicast Addresses

5.2. IPv6 Neighbor Discovery

In IPv6 the Neighbor Discovery (ND) protocol is seen as a replacement for IPv4 ARP message. It uses ICMPv6 messages with various purposes include finding a link-layer address of a neighbor, discover neighbor routers, determine any change in the link-layer address, determine when a neighbor is down, send network information from router to hosts, which includes hop limit, MTU size, determining the network prefix used on a link, address auto configuration, and the default route information.

There many types of ICMPv6 messages among which five types of messages are used by the ND protocol. The five types of ICMPv6 messages are briefly described in the following section:

- Router Solicitation (ICMPv6 type 133): This message is sent by hosts to routers as a request to router advertisement message. It uses a destination multicast address: FF02::2
- Router Advertisement Messages (ICMPv6 type 134): This message is used by routers to announce its presence in a network. The message includes network information related to IPv6 prefixes, default route, MTU size, hop limit and auto configuration flag. It uses a destination multicast address: FF02::1
- Neighbor Solicitation Messages (ICMPv6 type 135): This message is sent by hosts to determine the existence of another host on the same. The goal is to find the link-layer of neighbor nodes on the same link.
- Neighbor Advertisement Messages (ICMPv6 type 136): This message is sent by hosts to indicate the existence of the host and it provides information about its own link-layer address.
- Redirect Messages (ICMPv6 type 137): This message is sent by a router to inform a host about a better router to reach a particular destination address.

In RX1000, Neighbor Discovery should be configured on all Ethernet interfaces enabled for IPv6. The following figure displays the available configuration options for IPv6 Neighbor Discovery.

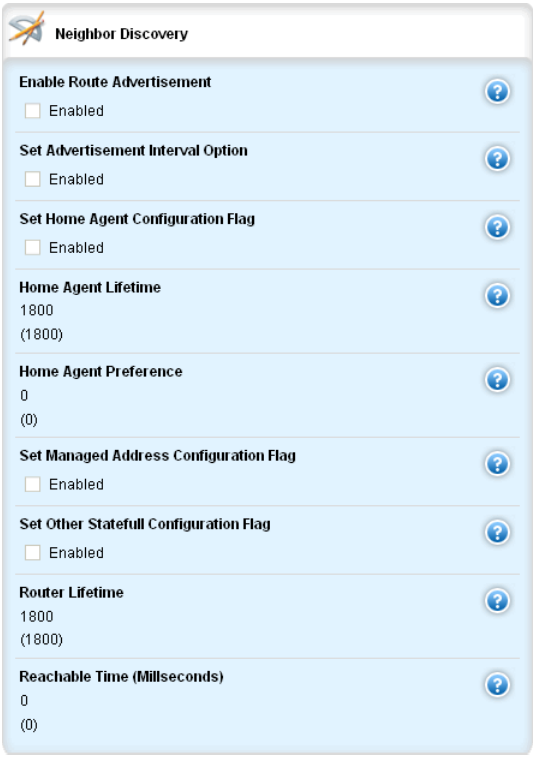

Figure 5.1. Neighbor Discovery form

The path to the Neighbor Discovery form is *ip/{interface}/ipv6/nd*.

Enable Route Advertisement

Enable to send router advertisement messages.

Set Advertisement Interval Option

Includes an Advertisement Interval option which indicates to hosts the maximum time in milliseconds, between successive unsolicited router advertisements.

Set Home Agent Configuration Flag

Sets/unsets the flag in IPv6 router advertisements which indicates to hosts that the router acts as a home agent and includes a home agent option.

Home Agent Lifetime

Synopsis: *unsigned integer*

Default: *1800*

The value to be placed in the home agent option, when the home agent config flag is set, which indicated the home agent lifetime to hosts. A value of 0 means to place a router lifetime value.

Home Agent Preference

Synopsis: *unsigned integer*

Default:

The value to be placed in the home agent option, when the home agent config flag is set, which indicates the home agent preference to hosts.

Set Managed Address Configuration Flag

The flag in IPv6 router advertisements, which indicates to hosts that they should use the managed (stateful) protocol for addresses autoconfiguraiton in addition to any addresses autoconfigured using stateless address autoconfiguration.

Set Other Statefull Configuration Flag

The flag in IPv6 router advertisements, which indicates to hosts that they should use the administered (stateful) protocol to obtain autoconfiguration information other than addresses.

Router Lifetime

Synopsis: *unsigned integer*

Default: *1800*

The value (in seconds) to be placed in the Router Lifetime field of router advertisements sent from the interface. Indicates the usefulness of the router as a default router on this interface. Setting the value to zero indicates that the router should not be considered a default router on this interface. It must be either zero or between the value specified with the IPv6 nd ra-interval (or default) and 9000 seconds. The default is 1800 seconds.

Reachable Time (Millseconds)

Synopsis: *unsigned integer*

Default:

The value (in milliseconds) to be placed in the Reachable Time field in the router advertisement messages sent by the router. The configured time enables the router to detect unavailable neightbors. The value zero means unspecified (by this router). The default is 0.

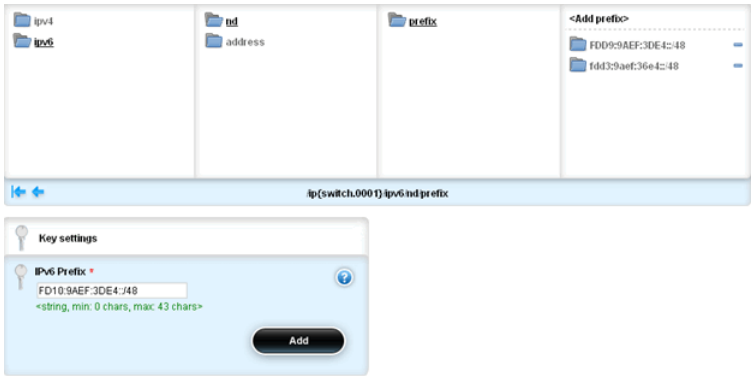

Figure 5.2. Neighbor Discovery IPv6 Prefix

An IPv6-capable interface can use Neighbor Discovery to advertise IPv6 network prefixes to its neighbor on the same link.

Figure 5.3. Neighbor Discovery IPv6 Prefix forms

IPv6 Prefix

Synopsis: *IPv6 address and prefix in CIDR notation*

The IPv6 network/prefix.

Valid Lifetime

Synopsis: *unsigned integer*

Synopsis: *string - the keyword { infinite }*

The length of time in seconds during what the prefix is valid for the purpose of on-link determination. The default value is 2592000.

Preferred Lifetime

Synopsis: *unsigned integer*

Synopsis: *string - the keyword { infinite }*

The length of time in seconds during which addresses generated from the prefix remain preferred. The default value is 604800.

Off Link

Indicates that advertisement makes no statement about on-link or off-link properties of the prefix.

No Autoconfig

Indicates to hosts on the local link that the specified prefix cannot be used for IPv6 autoconfiguration.

Set Router Address Flag

Indicates to hosts on the local link that the specified prefix contains a complete IP address by setting the R flag.

This screen is accessible after adding an IPv6 Prefix under the Neighbor Discovery. To display the forms, navigate to *ip/{interface}/ipv6/nd/prefix*.

5.3. Non-switched Interface Menu

Figure 5.4. Non-switched Interface menu

The Non-switched (or Route-only) Interface menu is accessible from the main menu.

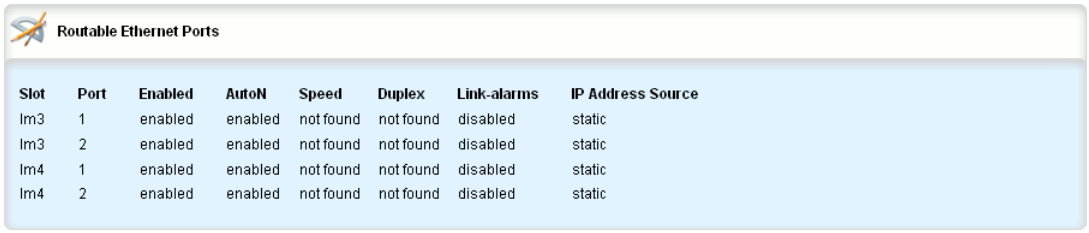
Routable Ethernet Ports
Slot Port Enabled AutoN Speed Duplex Link-alarms IP Address Source
lm3 1 enabled enabled not found not found disabled static
lm3 2 enabled enabled not found not found disabled static
lm4 1 enabled enabled not found not found disabled static
lm4 2 enabled enabled not found not found disabled static

Figure 5.5. Routable Ethernet Ports table

The path to the Routable Ethernet Ports table is *interface/eth*.

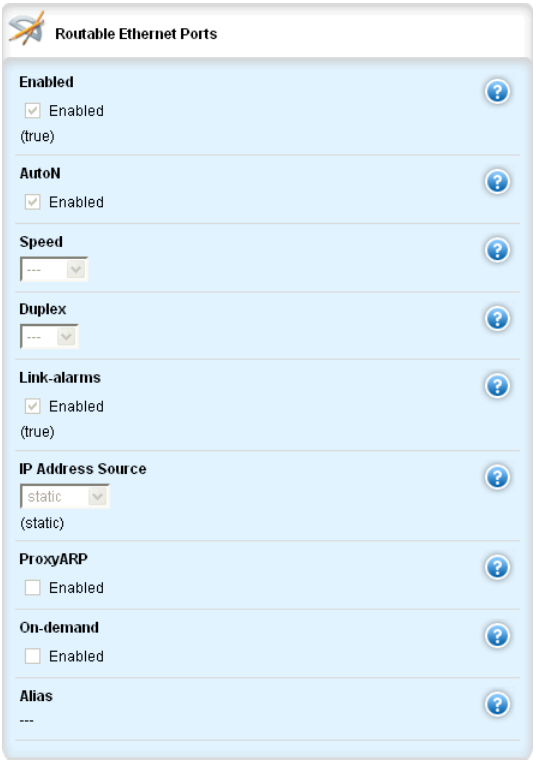

Figure 5.6. Routable Ethernet Ports form

The path to the Routable Ethernet Ports form is *interface/eth/{port}*.

Slot

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

Port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

Enabled

Synopsis: *boolean*

Default: *true*

Enables/Disables the network communications on this port

AutoN

Enables or disables IEEE 802.3 auto-negotiation. Enabling auto-negotiation results in speed and duplex being negotiated upon link detection; both end devices must be auto-negotiation compliant for the best possible results

Speed

Synopsis: *string - one of the following keywords { 1000, 100, 10 }*

Speed (in Megabit-per-second or Gigabit-per-second). If auto-negotiation is enabled, this is the speed capability advertised by the auto-negotiation process. If auto-negotiation is disabled, the port is explicitly forced to this speed mode. AUTO means advertise all supported speed modes.

Duplex

Synopsis: *string - one of the following keywords { full, half }*

If auto-negotiation is enabled, this is the duplex capability advertised by the auto-negotiation process. If auto-negotiation is disabled, the port is explicitly forced to this duplex mode. AUTO means advertise all supported duplex modes.

link-alarms

Synopsis: *boolean*

Default: *true*

Disabling link-alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that interface. Link alarms may also be controlled for the whole system under admin / alarm-cfg.

IP Address Source

Synopsis: *string - one of the following keywords { dynamic, static }*

Default: *static*

Whether the IP address is static or dynamically assigned via DHCP or BOOTP. Option DYNAMIC is a common case of dynamically assigned IP address. It switches between BOOTP and DHCP until it gets the response from the relevant server. Must be static for non-management interfaces

ProxyARP

Enables/Disables whether the port will respond to ARP requests for hosts other than itself

on-demand

This interface is up or down on demand of link fail over.

alias

Synopsis: *A string*

The SNMP alias name of the interface

5.3.1. Configuring IP Address Source and ProxyARP for Non-switched Interfaces

IP addresses on routable ports are static by default. To change the IP address of the port to dynamic, follow the procedure below. ProxyARP can also be enabled using this form.

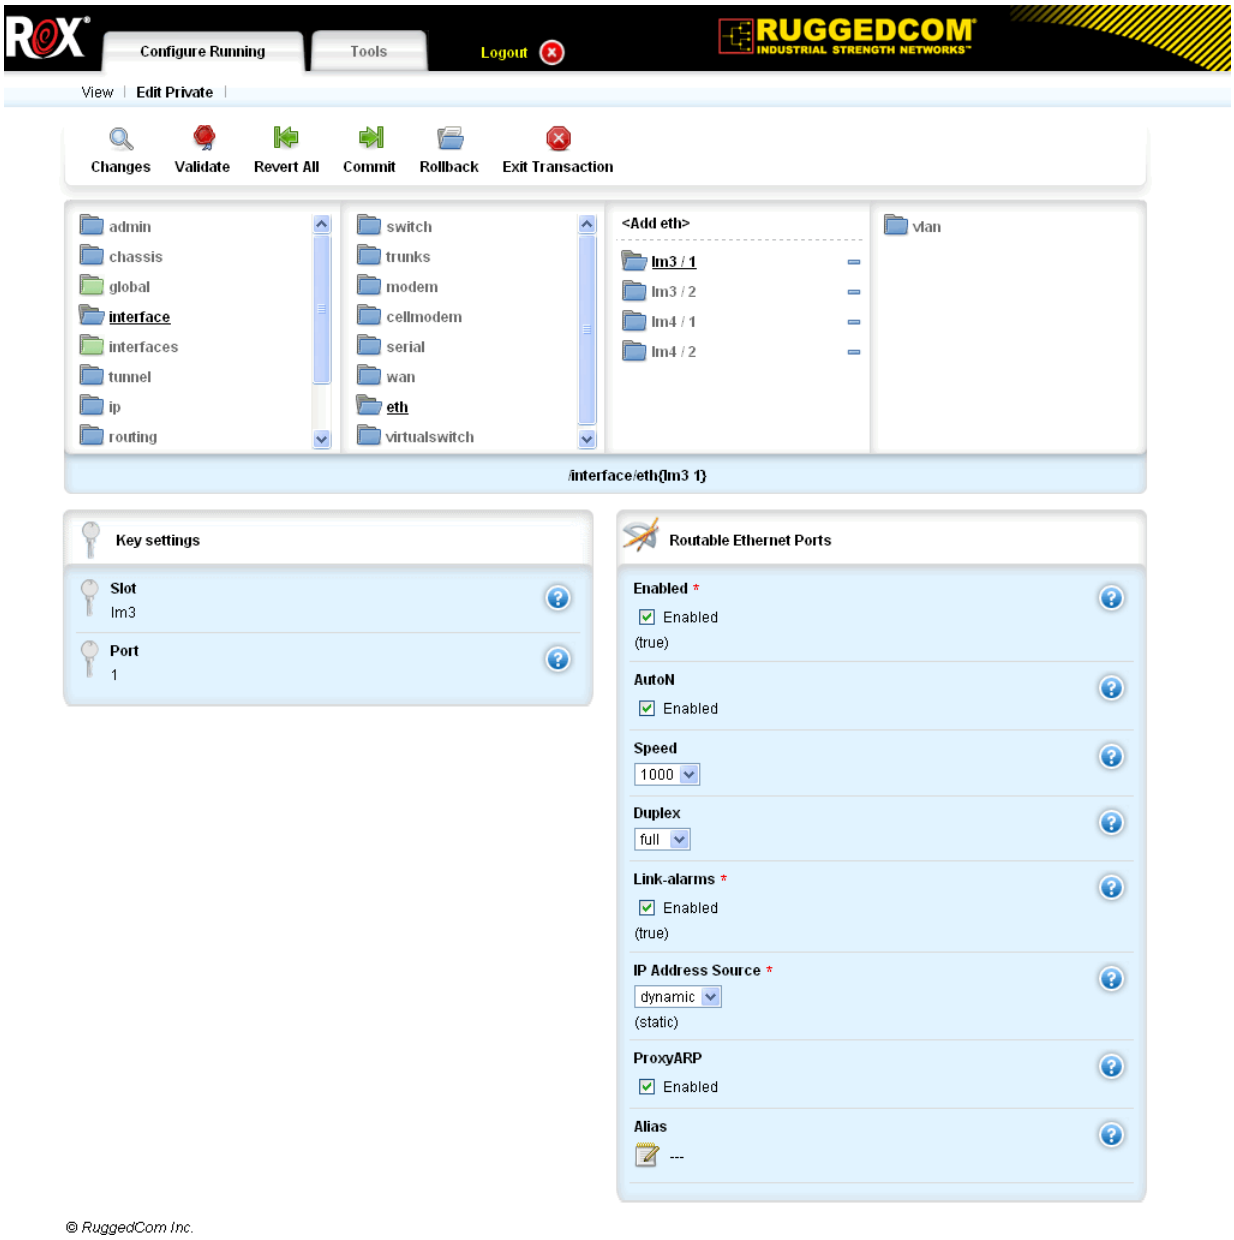

© RuggedCom Inc.

Figure 5.7. Configuring Dynamic Address Source and ProxyARP

Procedure 5.1. Configuring IP Address Source and ProxyARP for Non-switched Interfaces

1. Go into Edit Private mode.
2. Go to *interface/eth/{port}*. The Routable Ethernet Ports form appears.
3. In the *IP Address Source* field, select **dynamic** if you want the interface to get an IP address from a DHCP server. For information on configuring RX1000 as a DHCP server, see Chapter 14, *DHCP Server*. To assign a static IP address to an interface, see Chapter 4, *Basic Network Configuration*.
4. Click **Commit**.
5. Click **Exit Transaction**.

To set ProxyARP for a static or dynamic interface, follow the procedure below.

Procedure 5.2. Setting ProxyARP

1. Go into Edit Private mode.
2. Go to *interface/eth/(port}*. The Routable Ethernet Ports form appears.
3. In the *ProxyARP* field, click **Enabled**.
4. Click **Commit**.
5. Click **Exit Transaction**.

6. Alarms

6.1. Introduction

The ROXII alarm system is a highly configurable notification system of events of interest. Asserted alarms in the system may be viewed in a table in the CLI, web user interface, as well as queried by NETCONF. Alarms are categorized by subsystem.

The alarm system allows the user to:

- enable/disable alarms with the exception of mandatory alarms
- configure whether or not an alarm triggers the fail-relay and paints the alarm LED red
- configure the severity of an alarm to one of the following: emergency, alert, critical, error, warning, notice, info, debug (in descending order of severity). A small minority of alarms have fixed severity.

6.1.1. Alarm Subsystems

As of the current release, two subsystems support alarms: Admin and Chassis.

Note that some of the following examples describing the nature of each alarm subsystem may not be available in this release. A list of the available alarms can be viewed in the configuration file at **`/admin/alarm-cfg`**.

Admin Subsystem: these alarms are for administrative aspects of the device, including feature-key problems, upgrades, and configuration changes.

Chassis Subsystem: these alarms are for physical or electrical problems, or events of interest. This includes irregular voltages at the power supply or the insertion or removal of a module.

6.1.2. Fail-Relay Behavior

The fail-relay shall be activated when an active alarm in the system is also configured to trigger it. Once an alarm has been acknowledged or cleared it ceases to assert the fail-relay. The fail-relay will only be de-activated when all active alarms that are configured to assert it have been acknowledged or cleared.

6.1.3. Alarm LED Behavior

The alarm LED on the device faceplate is ON and red when unacknowledged alarm(s) are asserted and the LED is enabled for any of the active alarms. After an alarm is acknowledged or cleared, the LED is switched OFF.

6.1.4. Clearing and Acknowledging Alarms

There are two broad types of alarms:

1. Non-Clearable alarms - Users cannot clear these alarms, only acknowledge them; the difference between these actions is outlined later in this section. These alarms have a condition associated with them that the system assesses. The system asserts the alarm when the condition is true and clears the alarm when the condition has been resolved. An example of this is 'Bad input supply on power module'. If a redundant power module loses its supply an alarm is asserted. If the problem is resolved and power is returned to the module, the system de-asserts the alarm. De-asserted alarms remain as active alarms until acknowledged by the user.
2. Clearable alarms - these alarms simply report an event of interest that has no resolution per se. An example of this would be a 'configuration changed' alarm. These alarms are clearable by the user and are never cleared by the system.

Alarms may be cleared and acknowledged both on an individual basis and globally (i.e. clear/acknowledge all active-alarms). When an alarm is cleared by the user it is removed from the active

alarms table and no longer asserts the fail-relay and LED. When an alarm is acknowledged by the user it de-asserts the fail-relay and LED, but it remains in the active alarms table, unless the alarm is non-clearable and de-asserted by the system. In the latter case it is removed from the table, because the condition was resolved.

6.2. Alarm Configuration

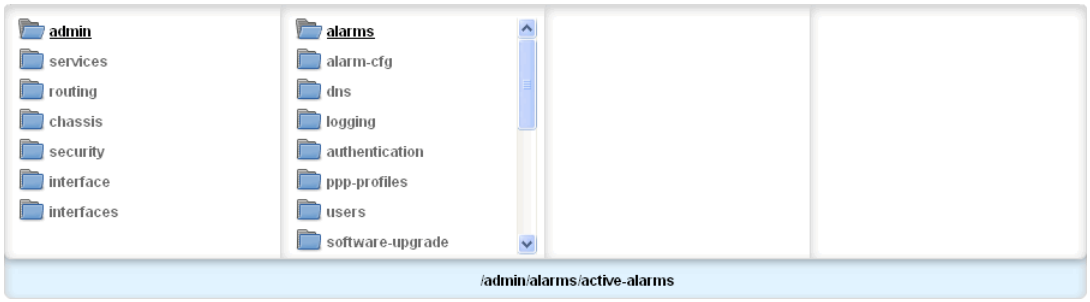

Figure 6.1. Alarms menu

The Alarms menu is accessible from the main menu under *admin*.

View active alarms in the Active Alarms table.

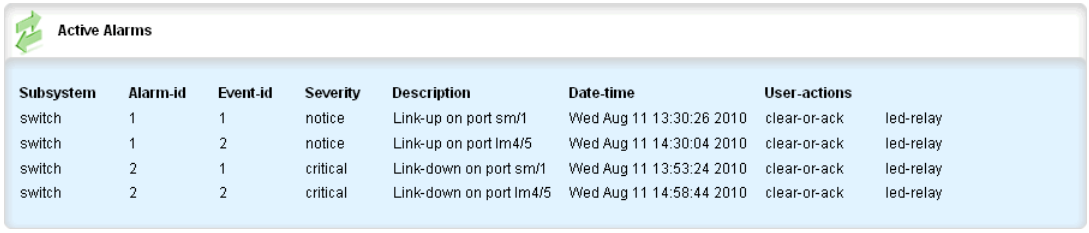
Active Alarms
Subsystem Alarm-id Event-id Severity Description Date-time User-actions
switch 1 1 notice Link-up on port sm/1 Wed Aug 11 13:30:26 2010 clear-or-ack led-relay
switch 1 2 notice Link-up on port lm4/5 Wed Aug 11 14:30:04 2010 clear-or-ack led-relay
switch 2 1 critical Link-down on port sm/1 Wed Aug 11 13:53:24 2010 clear-or-ack led-relay
switch 2 2 critical Link-down on port lm4/5 Wed Aug 11 14:58:44 2010 clear-or-ack led-relay

Figure 6.2. Active Alarms table

If data is configured, the Active Alarms table will appear on the same screen as the *Alarms* menu.

Figure 6.3. Active Alarms Key Settings form

If data is configured, the path to the Key Settings form and Active Alarms form is *admin/alarms/{interface}*.

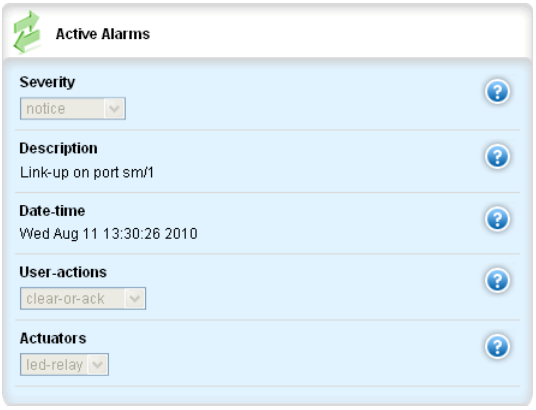

Figure 6.4. Active Alarms form

subsystem

Synopsis: *string - one of the following keywords { wan, switch, chassis, admin }*

Alarms are categorized by the subsystem to which they belong e.g.: Admin, Chassis, Ethernet, WAN.

Alarm ID

Synopsis: *integer*

Alarm Type Identifier. A value that uniquely defines a type of alarm.

Event ID

Synopsis: *integer*

Alarm Event Identifier. A value that uniquely defines a specific alarm event of the indicated alarm type.

severity

Synopsis: *string - one of the following keywords { debug, info, notice, warning, error, critical, alert, emergency }*

The class of severity: Emergency, Alert, Critical, Error, Warning, Notice, Info, Debug

description

Synopsis: *string*

When applicable, provides further details on the alarmable event

Date/Time

Synopsis: *string*

The date and time the event was detected

User Actions

Synopsis: *string - one of the following keywords { must-resolve, clear-or-ack, resolve-or-ack }*

There are three categories of alarms:

1. clear or ack : the user can clear (remove from 'active-alarm' list) and/or acknowledge (turn off actuator(s) but keep as active-alarm).
2. ack or resolve : the user can acknowledge only, the system will clear the alarm once it is acknowledged and the condition is resovled.
3. must-resolve : for a minority of alarms, the condition must be resolved to turn off actuators and clear the alarm.

actuators

Synopsis: *string - one of the following keywords { acked, none, led-relay, led, relay }*

Indicates which actuator(s) this alarm currently asserts. 'ACKED' indicates the alarm was acknowledged so actuators are de-asserted.

Individual alarms can be cleared or acknowledged on the Clear Alarm Menu Action form or the Acknowledge Alarm Menu Action form. To clear or acknowledge an alarm, select *admin/alarms/{alarms submenu}* and then select the *Clear* action or the *Acknowledge* action.

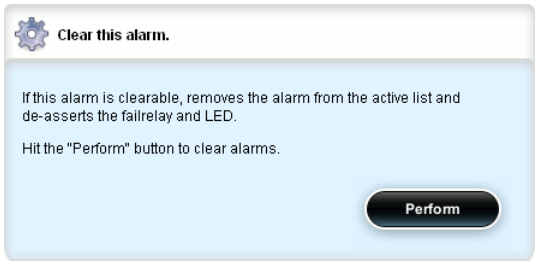
Clear this alarm.
If this alarm is clearable, removes the alarm from the active list and de-asserts the failrelay and LED.
Hit the "Perform" button to clear alarms.
Perform

Figure 6.5. Clear Alarm Menu Action form

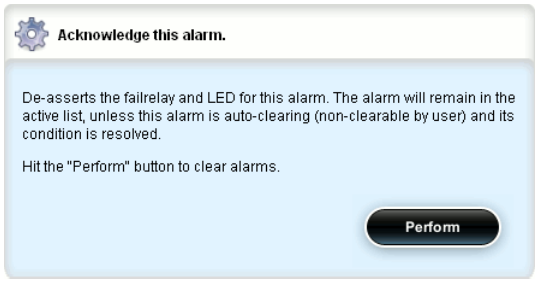
Acknowledge this alarm.
De-asserts the failrelay and LED for this alarm. The alarm will remain in the active list, unless this alarm is auto-clearing (non-clearable by user) and its condition is resolved.
Hit the "Perform" button to clear alarms.
Perform

Figure 6.6. Acknowledge Alarm Menu Action form

To clear or acknowledge ALL alarms, instead of only individual alarms, access the Clear All Alarms and Acknowledge All Alarms menu action forms. These forms are accessible from the *admin* menu. The path to the Clear All Alarms Menu Action and the Acknowledge All Alarm Menu Action is *admin*, then clicking on the *clear-all-alarms* action or the *acknowledge-all-alarms* action.

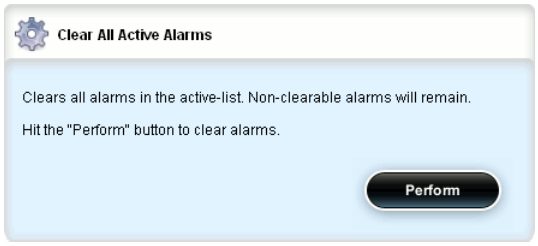
Clear All Active Alarms
Clears all alarms in the active-list. Non-clearable alarms will remain.
Hit the "Perform" button to clear alarms.
Perform

Figure 6.7. Clear All Alarms Menu Action form

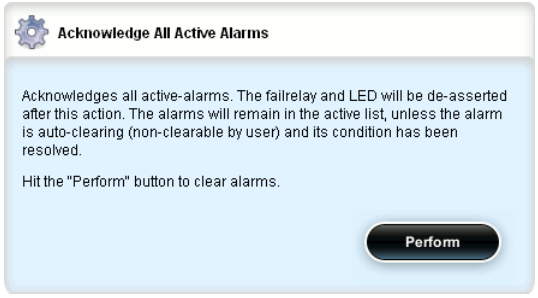
Acknowledge All Active Alarms
Acknowledges all active-alarms. The failrelay and LED will be de-asserted after this action. The alarms will remain in the active list, unless the alarm is auto-clearing (non-clearable by user) and its condition has been resolved.
Hit the "Perform" button to clear alarms.
Perform

Figure 6.8. Acknowledge All Alarms Menu Action form

6.2.1. Administrative Alarm Configuration

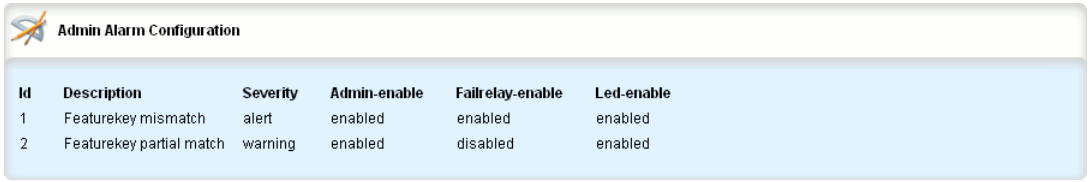
Admin Alarm Configuration
Id Description Severity Admin-enable Failrelay-enable Led-enable
1 Featurekey mismatch alert enabled enabled enabled
2 Featurekey partial match warning enabled disabled enabled

Figure 6.9. Admin Alarm Configuration table

The path to the Admin Alarm Configuration table is *admin/alarm-config/admin*.

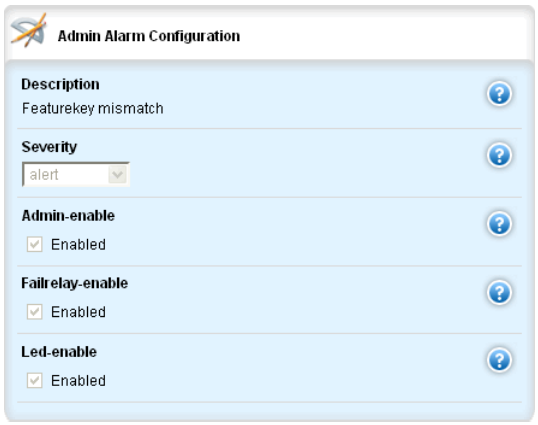

Figure 6.10. Admin Alarm Configuration form

The path to the Admin Alarm Configuration form is *admin/alarm-config/admin/{alarm id}*.

id

Synopsis: *integer*

This is the ID number of the alarm assigned by the system.

description

Synopsis: *A string*

The name of the alarm.

severity

Synopsis: *string - one of the following keywords { debug, info, notice, warning, error, critical, alert, emergency }*

The severity level can be one of emergency, alert, critical, error, warning, notice, info, and debug. This cannot be changed for some alarms

admin-enable

If disabled, the alarm is not reported in the active list and does not actuate led/failrelay.

failrelay-enable

If enabled, this alarm will assert the failrelay.

led-enable

If enabled, the main 'Alarm' LED light will be red when this alarm is asserted. If disabled, the main 'Alarm' LED light is not affected by this alarm.

6.2.2. Chassis Alarm Configuration

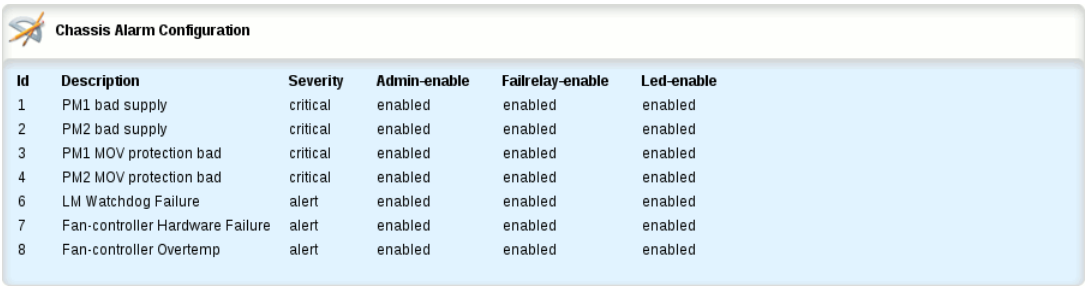

Id	Description	Severity	Admin-enable	Failrelay-enable	Led-enable
1	PM1 bad supply	critical	enabled	enabled	enabled
2	PM2 bad supply	critical	enabled	enabled	enabled
3	PM1 MOV protection bad	critical	enabled	enabled	enabled
4	PM2 MOV protection bad	critical	enabled	enabled	enabled
6	LM Watchdog Failure	alert	enabled	enabled	enabled
7	Fan-controller Hardware Failure	alert	enabled	enabled	enabled
8	Fan-controller Overtemp	alert	enabled	enabled	enabled

Figure 6.11. Chassis Alarm Configuration table

The path to the Chassis Alarm Configuration form is *admin/alarm-config/chassis*.

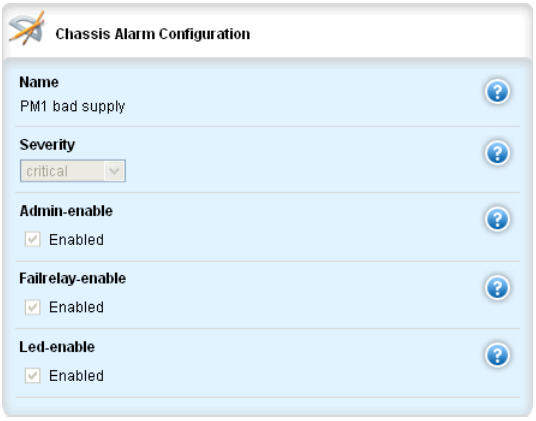

Figure 6.12. Chassis Alarm Configuration form

The path to the Chassis Alarm Configuration form is *admin/alarm-config/chassis/{alarm id)*.

id

Synopsis: *integer*

This is the ID number of the alarm assigned by the system.

description

Synopsis: *A string*

The name of the alarm.

severity

Synopsis: *string - one of the following keywords { debug, info, notice, warning, error, critical, alert, emergency }*

The severity level can be one of emergency, alert, critical, error, warning, notice, info, and debug. This cannot be changed for some alarms

admin-enable

If disabled, the alarm is not reported in the active list and does not actuate led/failrelay.

failrelay-enable

If enabled, this alarm will assert the failrelay.

led-enable

If enabled, the main 'Alarm' LED light will be red when this alarm is asserted. If disabled, the main 'Alarm' LED light is not affected by this alarm.

7. Domain Name Search

7.1. Domain Name Lookup

The DNS (Domain Name Service) menu is accessible from the main menu under *admin*. The path to this menu is *admin/dns*.

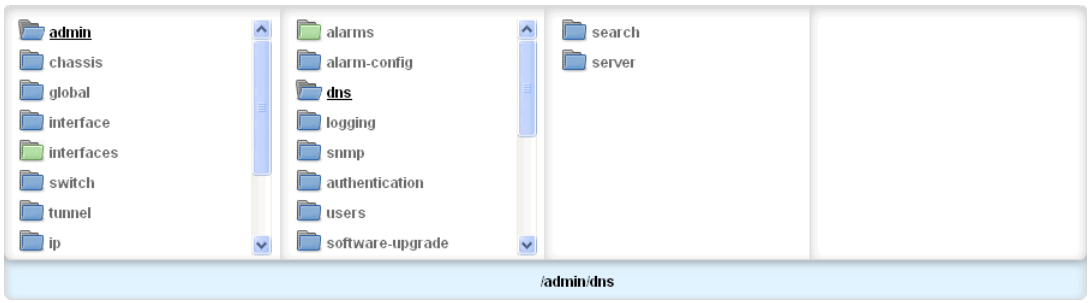

Figure 7.1. DNS menu

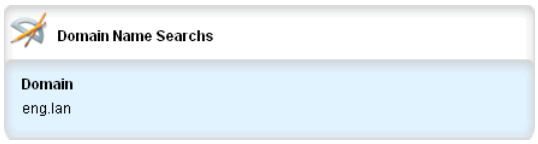

Figure 7.2. Domain Name Searches form

The path to the Domain Name Searches form is *admin/dns/search*.

domain

Synopsis: *Domain name (RFC 1034)*

Figure 7.3. Domain Name Servers

The path to the Domain Name Servers table is *admin/dns/server*.

address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

8. Logging

The syslog provides users with the ability to configure local and remote syslog connections. The remote syslog protocol, defined in RFC 3164, is a UDP/IP-based transport that enables a device to send event notification messages across IP networks to event message collectors, also known as syslog servers. The protocol is simply designed to transport these event messages from the generating device to the collector.

ROX™ supports up to 5 collectors (syslog servers). Remote Syslog provides the ability to configure:

- IP address(es) of collector(s).
- Source UDP port.
- Destination UDP port per collector.
- Syslog source facility ID per collector (same value for all ROX™ modules).
- Filtering severity level per collector (in case different collectors are interested in syslog reports with different severity levels).

8.1. Configuring Local Syslog

The local syslog configuration enables users to control what level of syslog information will be logged. Only messages of a severity level equal to or greater than the configured severity level are written to the syslog.txt file in the unit.

8.2. Configuring the Remote Syslog Server

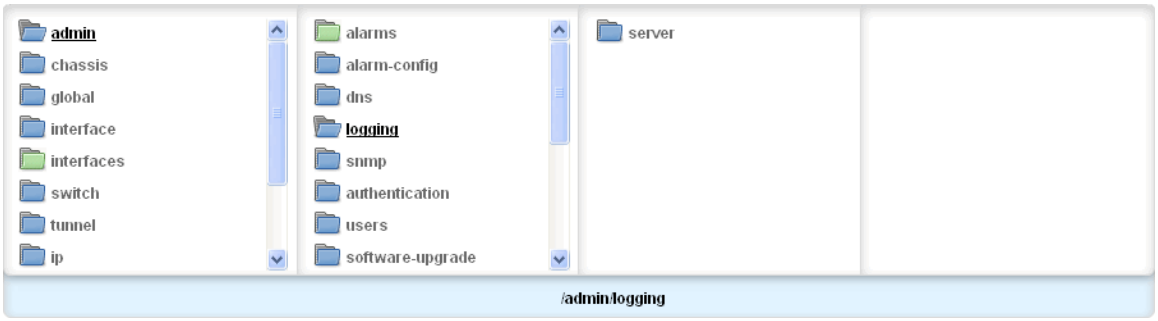

Figure 8.1. Logging menu

The Logging menu is accessible from the main menu under *admin*. The path to this menu is *admin/logging*.

Figure 8.2. Remote Server table

The Remote Server table appears on the same screen as the *Logging* menu.

The Remote Server table can be used to identify a remote logging server.

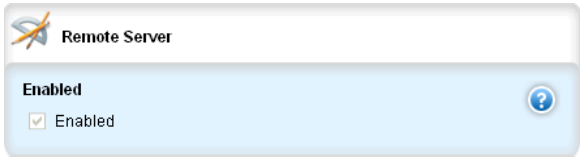

Figure 8.3. Remote Server form

If data is configured, there will be a list of logging servers under *admin/logging/server*. Clicking on each server will allow you to access the settings and Remote Server forms.

Server IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Synopsis: *Domain name (RFC 1034)*

The IPv4 or IPv6 address of a logging server. Up to 8 logging servers can be added.

enabled

Enables/disables the feed to the remote logging server

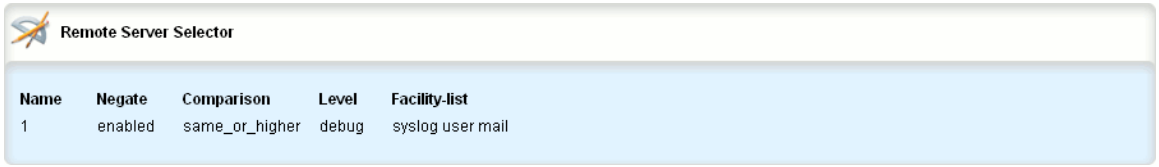

Figure 8.4. Remote Server Selector table

If data is configured, the path to the Remote Server Selector table will be *admin/logging/server*.

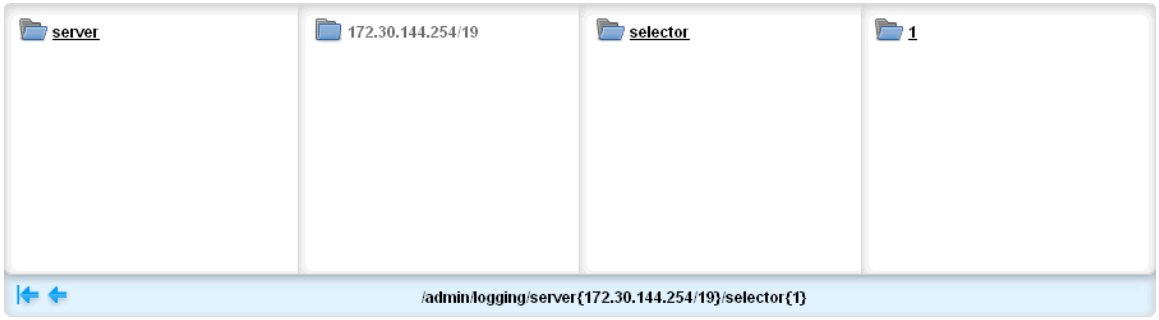

Figure 8.5. Selector menu

If data is configured, the path to the Remote Server Selector Forms (below) will be *admin/logging/server*. Then click on the next linked submenu, then on "selector" and then "1" or any linked submenus that may be in this list.

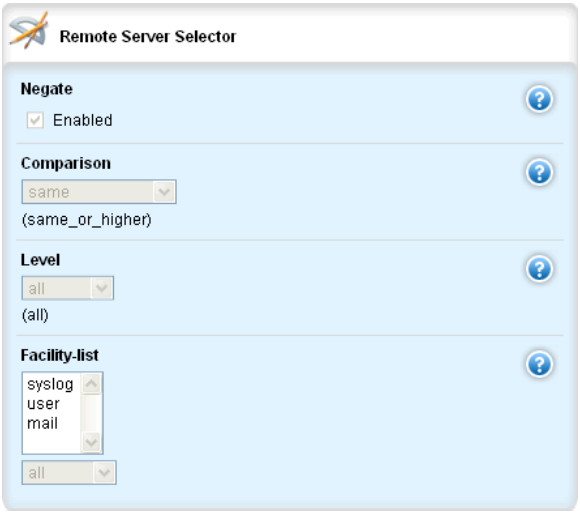

Figure 8.6. Remote Server Selector form

name

Synopsis: *integer*

The log selector identifier. Enter an integer greater than 0; up to 8 selectors can be added. The log selector determines which subsystem messages are included in the log.

negate

Excludes messages defined in the *Remote Server Selector* fields from the log. Selecting this option acts as a logical NOT for the selector definition.

For example: Selecting **same**, **debug**, and **mail** in the *Comparison*, *Level*, and *Facility-list* fields includes debug messages from the mail subsystem in the log. Selecting **Negate** *excludes* debug messages from the mail subsystem from the log.

comparison

Synopsis: *string - one of the following keywords { same, same_or_higher }*

Default: *same_or_higher*

The message severity levels to include in the log:

- **same:** includes only messages of the severity level selected in the *Level* field.
- **same_or_higher:** includes messages of the severity level selected in the *Level* field, and all messages of higher severity.

For example:

- Selecting **debug** in the *Level* field and **same** in the *Comparison* field includes only debug messages in the log.
- Selecting **debug** in the *Level* field and **same_or_higher** in the *Comparison* field includes debug and all higher severity messages in the log.

level

Synopsis: *string - one of the following keywords { all, none, debug, info, notice, warning, err, crit, alert, emerg }*

Default: *all*

The base message severity level to include in the log. **all** includes all messages. **none** excludes all messages. Other levels are listed in order of increasing severity.

facility-list

Synopsis: *string - one of the following keywords { all, local7, local6, local5, local4, local3, local2, local1, local0, uucp, user, syslog, security, news, mail, lpr, kern, ftp, daemon, cron, authpriv, auth }*

Synopsis: *"facility-list" occurs in an array of at most 8 elements.*

The subsystems generating log messages. Messages from the selected subusystems are included in the log. At least one subsystem must be selected; up to 8 subsystems can be selected.

8.3. Deleting Logs

For information on how to delete log files, see Section 2.10, “Deleting Log Files”.

9. SNMP

The SNMP (the Simple Network Management Protocol) protocol is used by network management systems and the devices they manage. SNMP is used to manage items on the device to be managed, as well as by the device itself, to report alarm conditions and other events.

The first version of SNMP, V1, provides the ability to send a notification of an event via "traps". Traps are unacknowledged UDP messages and may be lost in transit. SNMP V2 adds the ability to notify via "informs". Informs simply add acknowledgment to the trap process, resending the trap if it is not acknowledged in a timely fashion.

SNMP V1 and V2 transmit information in clear text (which may or may not be an issue depending on the facilities the data is transmitted over) and are lacking in the ability to authenticate a user. SNMP V3 adds strong authentication and encryption.

ROX™ supports Simple Network Management Protocol Version 3 (SNMPv3). This protocol provides secure access to devices by a combination of authentication and encrypting packets over the network. The security features provided are:

- message integrity - ensuring that a packet has not been tampered with in-transit.
- authentication – determining the message is from a valid source.
- encryption – scrambling the contents of a packet to prevent it from being seen by an unauthorized source.

SNMPv3 provides security models and security levels. A security model is an authentication strategy that is set up for a user and the group in which the user resides. A security level is a permitted level of security within a security model. A combination of a security model and security level will determine which security mechanism is employed when handling an SNMP packet.

Note the following about SNMPv3 protocol:

- each user belongs to a group.
- a group defines the access policy for a set of users.
- an access policy defines what SNMP objects can be accessed for: reading, writing and creating notifications.
- a group determines the list of notifications its users can receive.
- a group also defines the security model and security level for its users.

9.1. SNMP Traps

The following SNMP traps are defined in the RX1000 MIB files:

Standard	MIB	Trap and Description
RFC 3418	SNMPv2-MIB	**authenticationFailure** An authenticationFailure trap signifies that the SNMP entity has received a protocol message that is not properly authenticated. While all implementations of SNMP entities MAY be capable of generating this trap, the snmpEnableAuthenTraps object indicates whether this trap will be generated.
		coldStart A coldStart trap signifies that the SNMP entity, supporting a notification originator application, is reinitializing itself and that its configuration may have been altered.
		warmStart A warmStart trap signifies that the SNMP entity, supporting a notification originator application, is reinitializing itself such that its configuration is unaltered.

Standard	MIB	Trap and Description
RFC 4188	BRIDGE-MIB	**newRoot** The newRoot trap indicates that the sending agent has become the new root of the Spanning Tree; the trap is sent by a bridge soon after its election as the new root, e.g., upon expiration of the Topology Change Timer, immediately subsequent to its election. Implementation of this trap is optional.
		topologyChange A topologyChange trap is sent by a bridge when any of its configured ports transitions from the Learning state to the Forwarding state, or from the Forwarding state to the Blocking state. The trap is not sent if a newRoot trap is sent for the same transition. Implementation of this trap is optional.
IEEE Std 802.1AB-2005	LLDP-MIB	**lldpRemTablesChange** An lldpRemTablesChange notification is sent when the value of lldpStatsRemTableLastChangeTime changes. It can be utilized by an NMS to trigger LLDP remote systems table maintenance polls. Note that transmission of lldpRemTablesChange notifications are throttled by the agent, as specified by the 'lldpNotificationInterval' object.
RFC 1229, 2863, 2233, 1573	IF-MIB	**linkUp** A linkUp trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links left the down state and transitioned into some other state (but not into the notPresent state). This other state is indicated by the included value of ifOperStatus.
		linkDown A linkDown trap signifies that the SNMP entity, acting in an agent role, has detected that the ifOperStatus object for one of its communication links is about to enter the down state from some other state (but not from the notPresent state). This other state is indicated by the included value of ifOperStatus.
RuggedCom	RUGGEDCOM-TRAPS-MIB	**trapGenericTrap** Used for "User Authentication Events" only. The main subtree for RuggedCom generic traps.
		trapPowerSupplyTrap The main subtree for RuggedCom power supply trap.
		trapSwUpgradeTrap The main subtree for RuggedCom software uppgrade trap.
		trapCfgChangeTrap The main subtree for RuggedCom configuration change trap.
		trapFanBankTrap The main subtree for RuggedCom fan bank trap.
		trapHotswapModuleStateChangeTrap The main subtree for RuggedCom fan hotswap module state change trap.

Table 9.1. SNMP Traps

9.2. SNMP Access Configuration

To configure SNMP access to ROX™, follow the procedures outlined in the example below.

9.2.1. Add an SNMP User ID

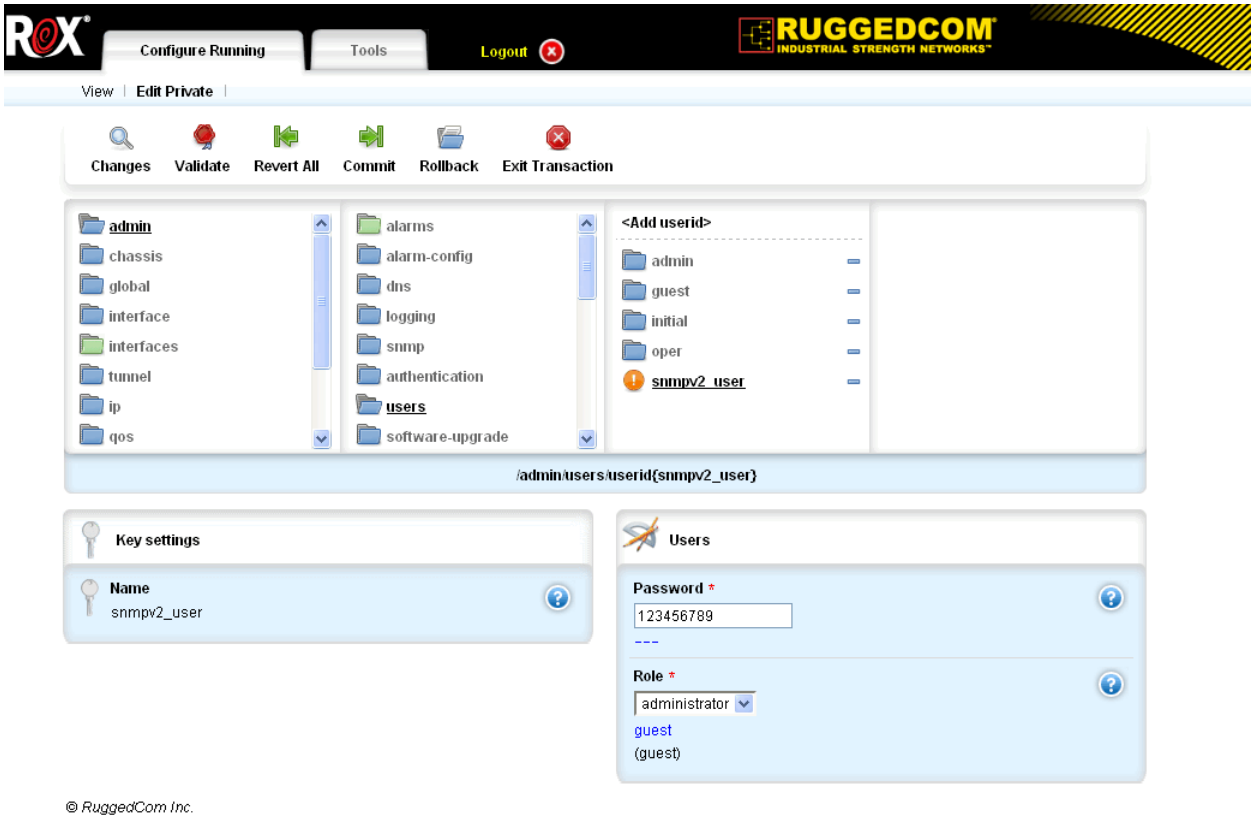

Figure 9.1. Adding an SNMP User ID

Procedure 9.1. Adding an SNMP User ID

1. Navigate to *admin/user*.
2. Click on <Add userid>. The Key settings form appears.
3. In the *Name* field, enter **snmpv2_user** and click **Add** . The Users form appears.
4. In the *Password* field, enter **123456789**.
5. In the *Role* field, enter **administrator**.
6. Click **Commit**.

9.2.2. Create an SNMP Community

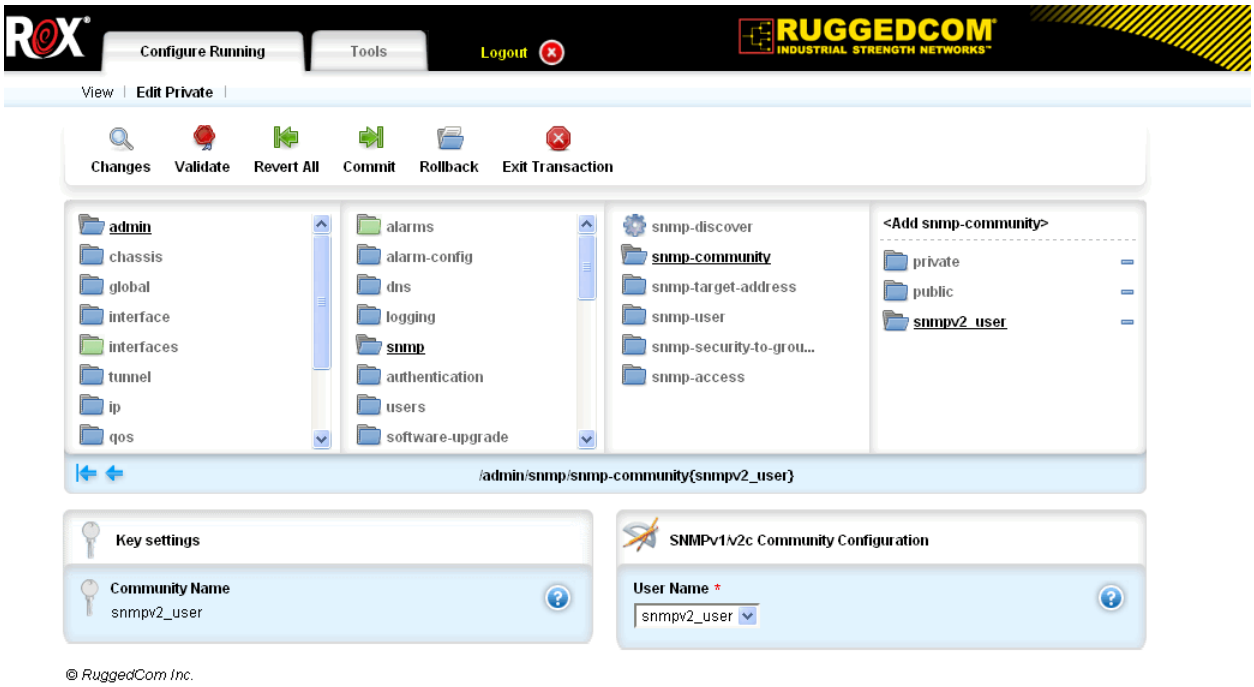

Figure 9.2. Creating an SNMP Community

Procedure 9.2. Creating an SNMP Community

1. Navigate to *admin/snmp/snmp-community*.
2. Click on <Add snmp-community>. The Key settings form appears.
3. In the *Community Name* field, enter **snmpv2_user** and click **Add**. The SNMPv1/v2c Community Configuration form appears.
4. In the *User Name* field, select **snmpv2_user**.
5. Click **Commit**.

9.2.3. Map the Community to a Security Group

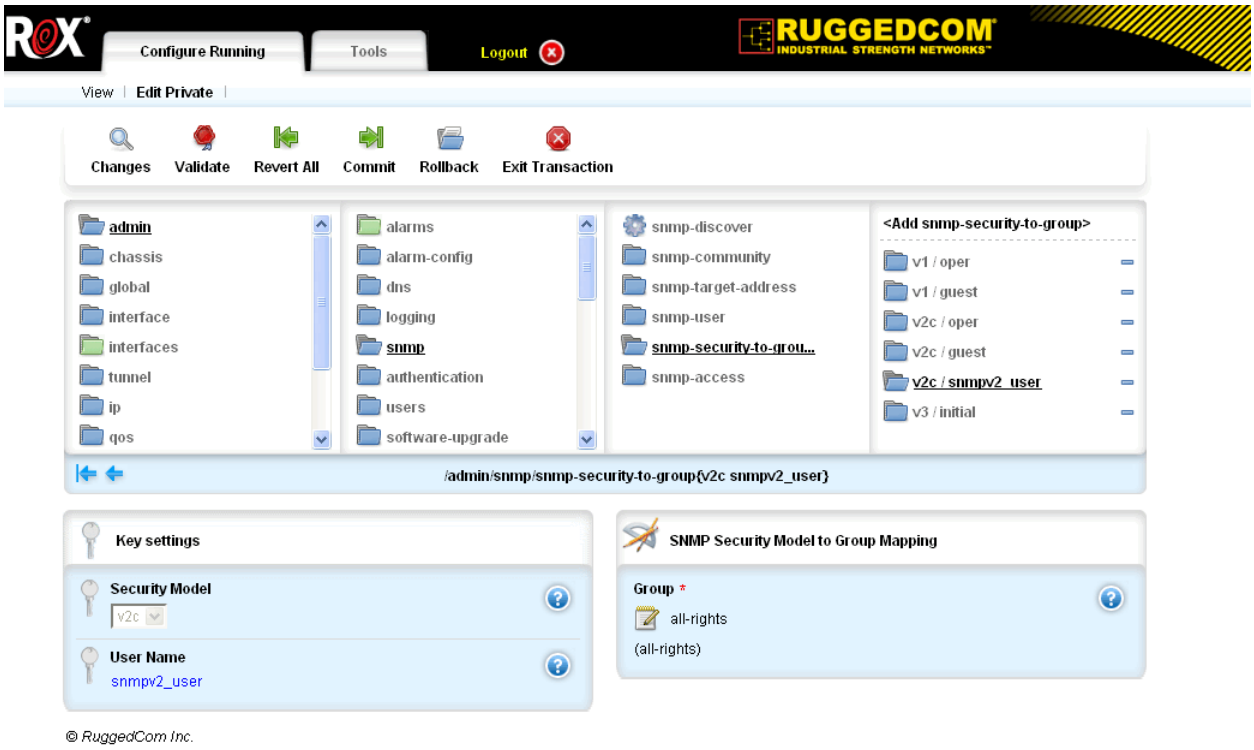

Figure 9.3. Mapping the Community to a Security Group

Procedure 9.3. Mapping the Community to a Security Group

1. Navigate to *admin/snmp/security-to-group*.
2. Click on <Add snmp-security-to-group>. The Key settings form appears.
3. In the *Security Model* field, select **v2c**.
4. In the *User Name* field, select **snmpv2_user** and click **Add**. The SNMP Security Model to Group Mapping form appears.
5. **all-rights** is the default in the Group field. Leave it as the default.
6. Click **Commit**.
7. Click **Exit Transaction**.

9.3. SNMP menu

Figure 9.4. SNMP menu

The SNMP menu is located at *admin/snmp*. The SNMP Sessions form and the SNMP USM Statistics form appear on the same screen as the *SNMP* menu.

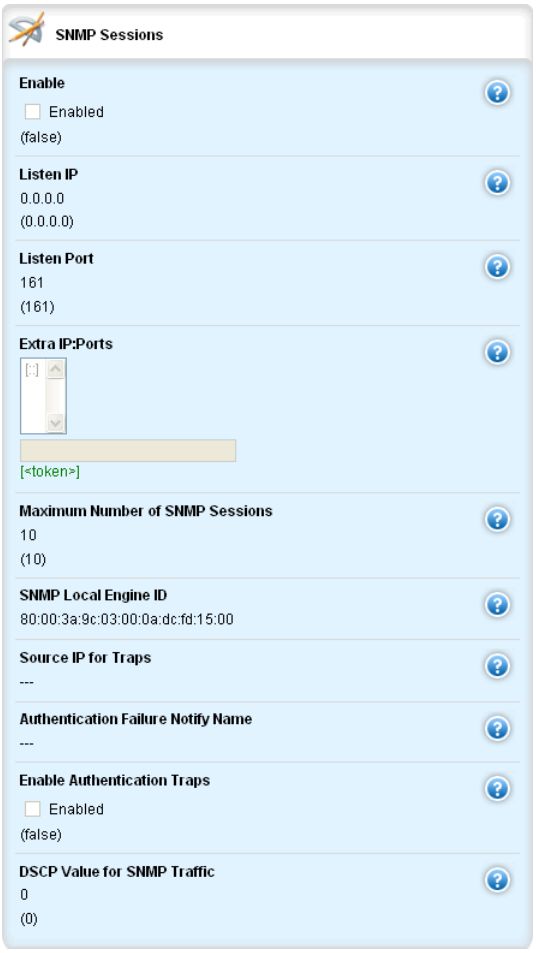

Figure 9.5. SNMP Sessions form

Enable

Synopsis: *boolean*

Default: *false*

Provides the ability to configure snmp features on the device.

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the SNMP agent will listen on for SNMP requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *161*

The port the SNMP agent will listen on for SNMP requests (default 161).

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

The SNMP agent will also listen on these IP Addresses:Port values. Add ':#' to set non-default port value #. (ie. xxx.xxx.xxx.xxx:19343 [::] [::]:16000)

Maximum Number of SNMP Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *30*

The maximum number of concurrent SNMP sessions

SNMP Local Engine ID

Synopsis: *A string*

Provides specific identification for the engine/device. By default, this value is set to use the base MAC address within the Engine ID value.

When using SNMP v3: if you change this value, you must also change the User SNMP Engine ID value for SNMP users.

Source IP for Traps

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

If set, all traffic/traps originating from this device shall use the configured IP Address for the Source IP

Authentication Failure Notify Name

Synopsis: *string - one of the following keywords { snmpv3_inform, snmpv3_trap, snmpv2_inform, snmpv2_trap, snmpv1_trap, none }*

Default: *none*

When the SNMP agent sends the standard authenticationFailure notification, it is delivered to the management targets defined for the snmpNotifyName in the snmpNotifyTable in SNMP-NOTIFICATION-MIB (RFC3413). If authenticationFailureNotifyName is the empty string (default), the notification is delivered to all management targets.

Enable Authentication Traps

Synopsis: *boolean*

Default: *false*

Enables authentication traps to be sent from the SNMP agent.

DSCP Value for SNMP Traffic

Synopsis: *unsigned byte integer in the range [0 to 63]*

Default:

Support for setting the Differentiated Services Code Point (6 bits) for traffic originating from the SNMP agent.

Figure 9.6. SNMP USM Statistics form

This table provides statistics for SNMP authentication requests

Unsupported Security Levels

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they requested a securityLevel that was unknown to the SNMP engine or otherwise unavailable.

Not In Time Windows

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they appeared outside of the authoritative SNMP engine's window.

Unknown User Names

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they referenced a user that was not known to the SNMP engine.

Unknown Engine IDs

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they referenced an snmpEngineID that was not known to the SNMP engine.

Wrong Digests

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they didn't contain the expected digest value.

Decryption Errors

Synopsis: *unsigned integer*

The total number of packets received by the SNMP engine which were dropped because they could not be decrypted.

9.4. SNMP Discovery

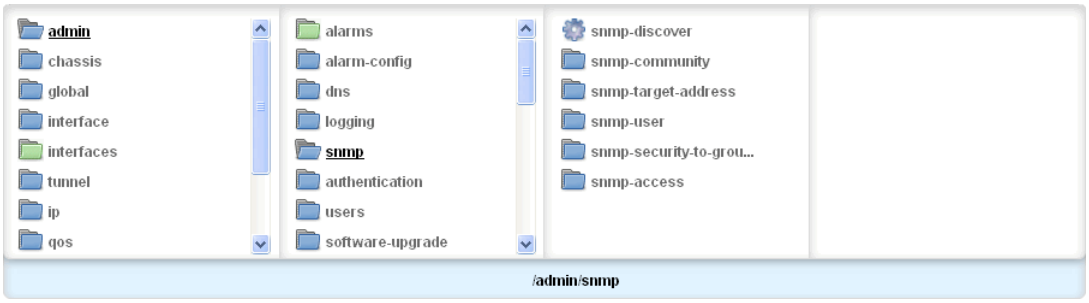

Figure 9.7. SNMP-Discover action

The path to this menu action is *admin/snmp/snmp-discover*.

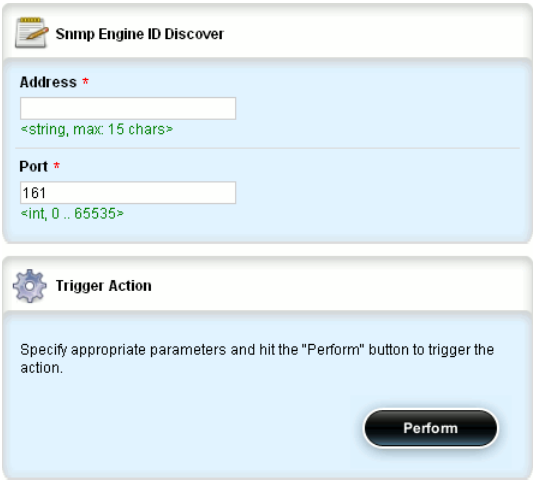

Figure 9.8. SNMP Engine ID Discover forms

To discover SNMP Engine IDs, use the SNMP Engine ID Discover and Trigger Action forms. On the *SNMP Engine ID Discover* form, enter parameters in the fields. On the *Trigger Action* form, click *Perform*.

9.5. SNMP Community

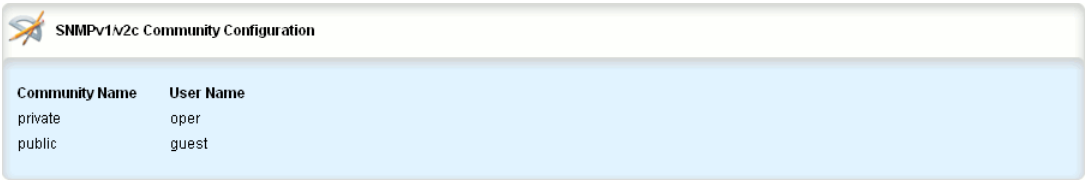

Figure 9.9. SNMPv1/v2c Community Configuration table

The path to the SNMP Community Configuration table is *admin/snmp/snmp-community*.

This table allows you to add and remove SNMPv1 and v2c Community Names.

Community Name

Synopsis: *A string*

The SNMP community name

User Name

Synopsis: *string*

The SNMP community security name

Figure 9.10. SNMPv1/v2c Community Configuration form

The path to the SNMP Community Configuration form is *admin/snmp/snmp-community/{private} or {public}*.

9.6. SNMP Target Addresses

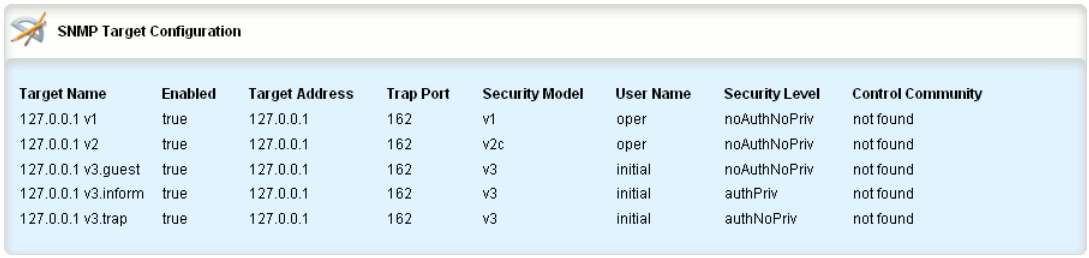

Target Name	Enabled	Target Address	Trap Port	Security Model	User Name	Security Level	Control Community
127.0.0.1 v1	true	127.0.0.1	162	v1	oper	noAuthNoPriv	not found
127.0.0.1 v2	true	127.0.0.1	162	v2c	oper	noAuthNoPriv	not found
127.0.0.1 v3.guest	true	127.0.0.1	162	v3	initial	noAuthNoPriv	not found
127.0.0.1 v3.inform	true	127.0.0.1	162	v3	initial	authPriv	not found
127.0.0.1 v3.trap	true	127.0.0.1	162	v3	initial	authNoPriv	not found

Figure 9.11. SNMP Target Configuration table

The path to the SNMP Target Configuration table is *admin/snmp/snmp-target-address*.

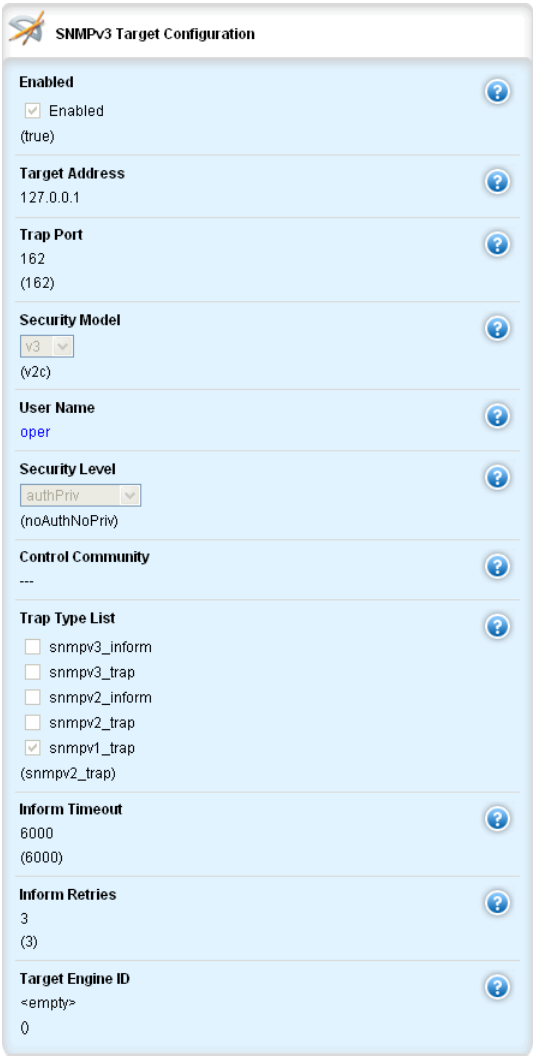

Figure 9.12. SNMPv3 Target Configuration form

To display the SNMP Target Configuration form, navigate to *admin/snmp/snmp-target-address/{address}*.

Target Name

A descriptive name for the target (ie. 'Corportate NMS')

enabled

Synopsis: *boolean*

Default: *true*

Enables/disables this specific target

Target Address

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

IPv4 or IPv6 address for the remote target.

Trap Port

Synopsis: *unsigned short integer*

Default: *162*

UDP Port for the remote target to receive traps on(default 162).

Security Model

Synopsis: *string - one of the following keywords { v3, v2c, v1 }*

Default: *v2c*

The SNMP security model to use: SNMPv1, SNMPv2c, or USM/SNMPv3

User Name

Synopsis: *string*

The user name to be used in communications with this target.

Security Level

Default: *noAuthNoPriv*

The SNMP security level:

- authPriv: communication with authentication and privacy.
- authNoPriv: communication with authentication and without privacy.
- noAuthnoPriv: communication without authentication and privacy.

Control Community

Synopsis: *A string*

Restrict incoming SNMP requests from the IPv4 or IPv6 address associated with this community.

Trap Type List

Default: *snmpv2_trap*

Selects the type of trap communications to be sent to this target

Inform Timeout

Default: *6000*

The timeout used for reliable inform transmissions (seconds*100)

Inform Retries

Default: *3*

The number of retries used for reliable inform transmissions

Target Engine ID

Default:

The target's SNMP local engine ID. This field may be left blank.

9.7. SNMP Users

These parameters provide the ability to configure users for the local SNMPv3 engine. Note that, if the security level employed is SNMPv1 or SNMPv2, User Name represents a community name for authentication or sending traps. Up to 32 entries can be configured.

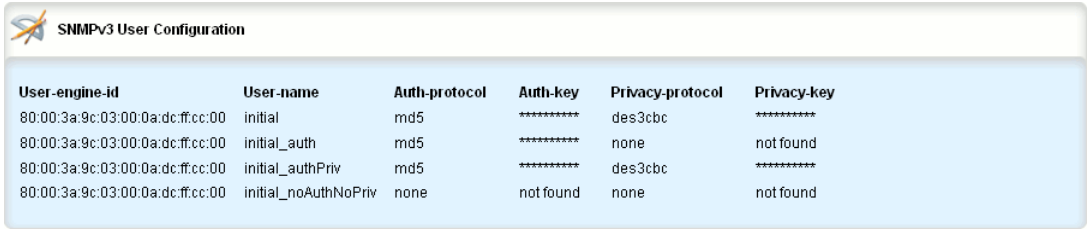

Figure 9.13. SNMP User Configuration table

The path to the SNMP User Configuration table is *admin/snmp/snmp-user*.

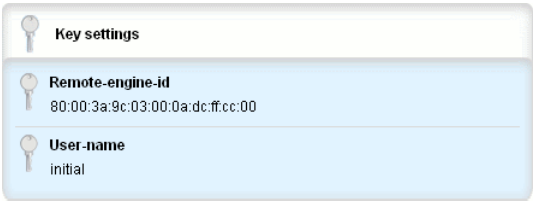

Figure 9.14. User Configuration Key Settings form

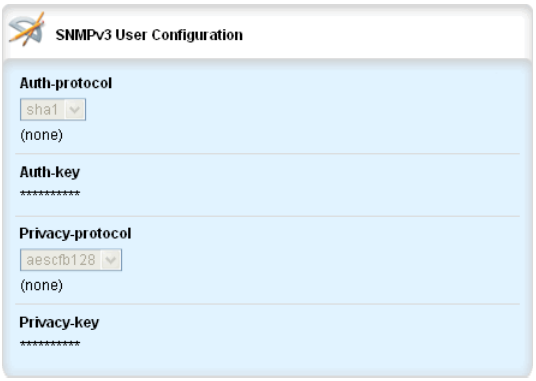

Figure 9.15. SNMP User Configuration form

The path to the Key Settings form and the SNMP User Configuration form is *admin/snmp/snmp-user/{user}*.

User SNMP Engine ID

The administratively-unique identifier for the SNMP engine; a value in the format nn:nn:nn:nn:nn:...:nn, where nn is a 2-digit hexadecimal number. Minimum length is 5 octets; maximum length is 32 octets. Each octet must be separated by a colon (:).

User Name

Synopsis: *string*

The user for the SNMP key. Select a user name from the list.

Authentication Protocol

Synopsis: *string - one of the following keywords { sha1, md5, none }*

Default: *none*

The authentication protocol providing data integrity and authentication for SNMP exchanges between the user and the SNMP engine.

Authentication Key

Synopsis: *"AES CFB128"-encrypted string*

A free-text password in the format **`$0$<your password>`**

Privacy Protocol

Synopsis: *string - one of the following keywords { aescfb128, des3cbc, none }*

Default: *none*

The symmetric privacy protocol providing data encryption and decryption for SNMP exchanges between the user and the SNMP engine.

Privacy Key

Synopsis: *"AES CFB128"-encrypted string*

A free-text password in the format **`$0$<your password>`**

9.8. SNMP Security to Group Maps

Entries in this table map the configuration of the security model and security name (user) into a group name, which is used to define an access control policy. Up to 32 entries can be configured.

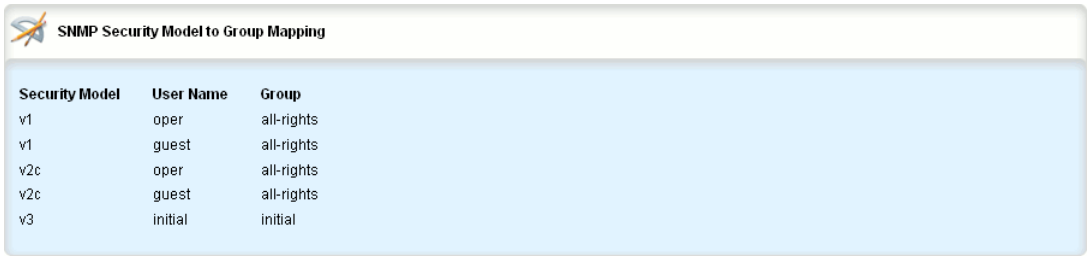

Figure 9.16. SNMP Security Model to Group Mapping table

The path to this table is *admin/snmp/snmp-security-to-group*.

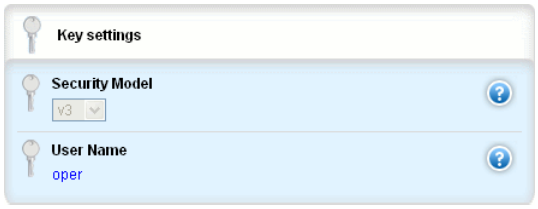

Figure 9.17. Key Settings form

Figure 9.18. SNMP Security Model to Group Mapping form

The path to these forms is *admin/snmp/snmp-security-to-group/{user}*.

Security Model

Synopsis: *string - one of the following keywords { v3, v2c, v1 }*

The SNMP security model to use: SNMPv1, SNMPv2c, or USM/SNMPv3

User Name

Synopsis: *string*

The security name (a ROX user name) for the SNMP group.

Group

Default: *all-rights*

The SNMP group name.

9.9. SNMP Access

These parameters provide the ability to configure access rights for groups. To determine whether access is allowed, one entry from this table needs to be selected and the proper view name from that entry must be used for access control checking. View names are predefined:

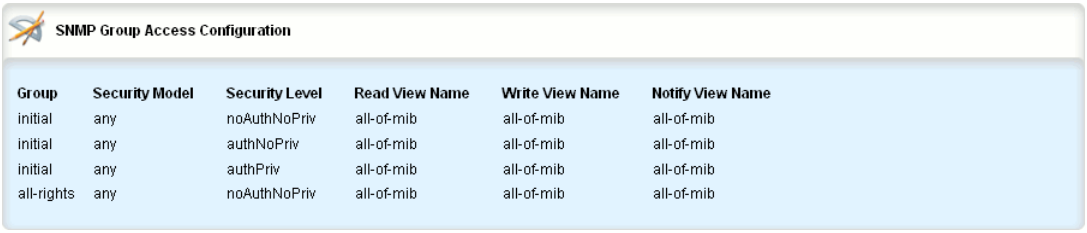
SNMP Group Access Configuration
Group Security Model Security Level Read View Name Write View Name Notify View Name
initial any noAuthNoPriv all-of-mib all-of-mib all-of-mib
initial any authNoPriv all-of-mib all-of-mib all-of-mib
initial any authPriv all-of-mib all-of-mib all-of-mib
all-rights any noAuthNoPriv all-of-mib all-of-mib all-of-mib

Figure 9.19. SNMP Group Access Configuration table

The path to this table is *admin/snmp/admin/snmp/snmp-access*.

Figure 9.20. Key Settings form

Figure 9.21. SNMP Group Access Configuration form

The path to this form is *admin/snmp/snmp-access/{access group}*.

Group

The SNMP group name.

Security Model

Synopsis: *string - one of the following keywords { v3, v2c, v1, any }*

The SNMP security model to use: SNMPv1, SNMPv2c, or USM/SNMPv3

Security Level

The SNMP security level:

- authPriv: communication with authentication and privacy.
- authNoPriv: communication with authentication and without privacy.
- noAuthnoPriv: communication without authentication and privacy.

Read View Name

Synopsis: *string - one of the following keywords { all-of-mib, restricted, v1-mib, no-view }*

Default: *all-of-mib*

The name of the read view to which the SNMP group has access: all-of-mib, restricted, v1-mib, or no-view.

Write View Name

Synopsis: *string - one of the following keywords { all-of-mib, restricted, v1-mib, no-view }*

Default: *all-of-mib*

The name of the write view to which the SNMP group has access: all-of-mib, restricted, v1-mib, or no-view.

Notify View Name

Synopsis: *string - one of the following keywords { all-of-mib, restricted, v1-mib, no-view }*

Default: *all-of-mib*

The name of the notification view to which the SNMP group has access: all-of-mib, restricted, v1-mib, or no-view.

10. Authentication

The Authentication menu is accessible from the main menu under *admin*. The path to this menu is *admin/authentication*.

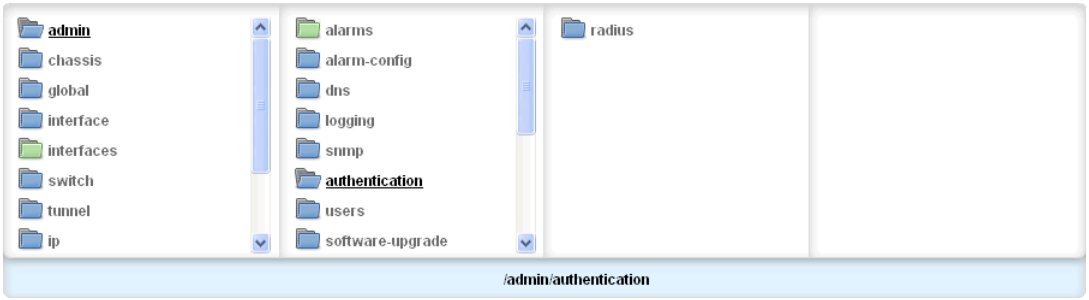

Figure 10.1. Authentication menu

The Authentication menu is accessible from the main menu under *admin*. The path to this menu is *admin/authentication*.

10.1. RADIUS

RADIUS (Remote Authentication Dial In User Service) is used to provide centralized authentication and authorization for network access. ROX™ assigns a privilege level of Admin, Operator or Guest to a user who presents a valid user name and password. The number of users who can access the ROX™ server is ordinarily dependent on the number of user records which can be configured on the server itself. ROX™ can also, however, be configured to pass along the credentials provided by the user to be remotely authenticated by a RADIUS server. In this way, a single RADIUS server can centrally store user data and provide authentication and authorization service to multiple ROX™ servers needing to authenticate connection attempts.

10.1.1. RADIUS overview

RADIUS (described in RFC 2865 [http://tools.ietf.org/html/rfc2865]) is a UDP-based protocol used for carrying authentication, authorization, and configuration information between a Network Access Server which desires to authenticate its links and a shared Authentication Server. RADIUS is also widely used in conjunction with 802.1x for port security using EAP (the Extensible Authentication Protocol, described in RFC 3748 [http://tools.ietf.org/html/rfc3748]).

A RADIUS server can act as a proxy client to other RADIUS servers or other kinds of authentication servers.

On receiving an authentication-authorization request from a client in an “Access-Request” packet, the RADIUS server checks the conditions configured for received username-password combination in the user database. If all the conditions are met, the list of configuration values for the user is placed into an “Access-Accept” packet. These values include the type of service (e.g. PPP, Login) and all the necessary values to deliver the desired service.

10.1.2. RADIUS Usage

The typical mode of operation involves a Network Access Server (NAS) - in this case the ROX™ - and a remote RADIUS server, where account information is stored. In the course of attempting to access connection-oriented services on the NAS, a user presents credentials to the NAS for authentication. The NAS forwards these to a configured RADIUS server and accepts from it the determination of whether the user is allowed the requested access. In order to protect the security of account information and of both the NAS and the RADIUS server, transactions are encrypted and authenticated through the use of a shared secret, which is never sent in the clear.

Some administrators set the passwords of existing ROX™ accounts uniquely for each router, and then employ a common password per account for all routers served by RADIUS. The router-specific passwords are restricted to a very few personnel. A larger set of expert users is granted the rights to SSH login using the RADIUS root account passwords.

10.1.3. RADIUS on ROX™

ROX™ supports RADIUS server redundancy. Multiple RADIUS servers, usually operating from a common database, may be used to authenticate a new session. If the first configured RADIUS server does not respond, subsequent servers will be tried until a positive/negative acknowledgment is received or an attempt has been made to contact all configured servers.

Each server is configured with an associated timeout which limits the time that ROX™ will wait for a response. An authentication request could thus require up to the sum of the timeouts of all configured servers.

RADIUS authentication activity is logged to the authorization log file, “auth.log”. Details of each authentication including the time of occurrence, source and result are included.

10.1.4. RADIUS, ROX™, and Services

RADIUS provides the means to restrict access on a per-service basis. Accounts may be configured on a RADIUS server to be allowed access only to the PPP service, for example. ROX™ supports RADIUS authentication for the following services:

- *LOGIN*
- *PPP*

ROX™ provides the option of designating different servers to authenticate LOGIN or PPP services separately or in combination.

The LOGIN Service

The *LOGIN* service consists of the following types of access:

- Local console logins via the serial port and modem
- Remote shell logins via SSH and Telnet
- Secure file transfers using SCP and SFTP (based on SSH)

Authentication requests for LOGIN services will attempt to use RADIUS first. If no response is received from any configured RADIUS server, ROX™ will authenticate against the local user database.

The PPP Service

The PPP service represents incoming PPP connections via modem. Authentication requests to the PPP service use RADIUS only. In the event that no response is received from any configured RADIUS server, ROX™ will not complete the authentication request.

10.1.5. RADIUS Authentication Configuration

There are two RADIUS server forms that can be configured in ROX™.

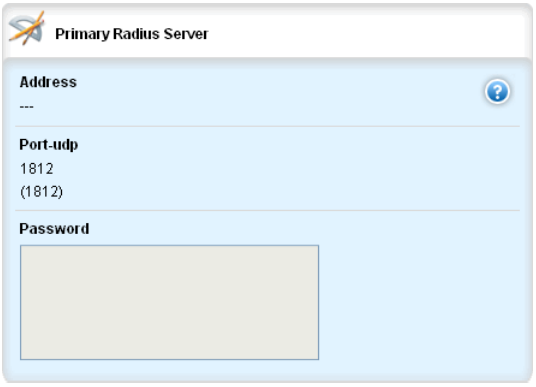

Figure 10.2. Primary RADIUS Server form

The Primary and Secondary RADIUS Server forms are accessible from the *radius* menu, which is a sub menu of the authentication menu. The path to this menu is *admin/authentication/radius*. These forms are also accessible from *global/ppp/radius*.

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

port-udp

Synopsis: *integer*

Default: *1812*

The network port of the server

password

Synopsis: *"AES CFB128"-encrypted string*

The password of the RADIUS server

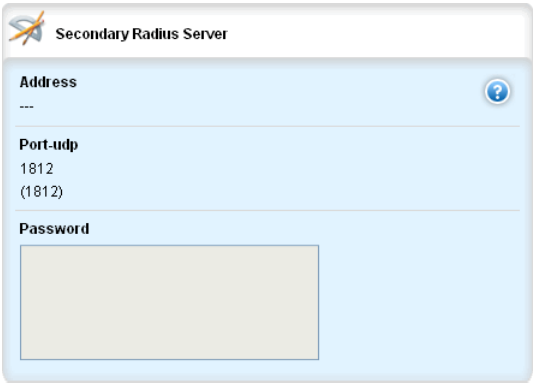

Figure 10.3. Secondary RADIUS Server form

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

port-udp

Synopsis: *integer*

Default: *1812*

The network port of the server

password

Synopsis: *"AES CFB128"-encrypted string*

The password of the RADIUS server

For additional information on RADIUS server configuration, please see Appendix B, *RADIUS Server Configuration*.

11. NETCONF

Figure 11.1. NETCONF menu

The NETCONF menu is accessible from the main menu under *admin*. The path to this menu is *admin/netconf*.

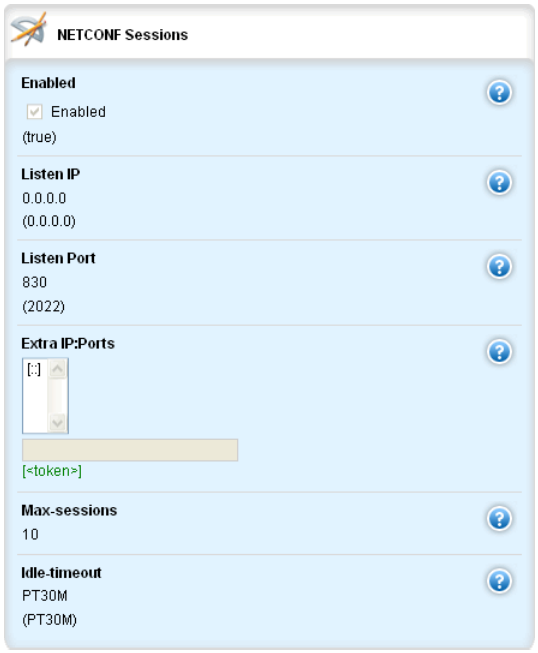

Figure 11.2. NETCONF Sessions form

The path to the NETCONF Sessions form and the NETCONF State/Statistics form is *admin/netconf*.

enabled

Synopsis: *boolean*

Default: *true*

Provides the ability to configure NETCONF features on the device.

Listen IP

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *IPv6 address in colon-separated hexadecimal notation*

Default: *0.0.0.0*

The IP Address the CLI will listen on for NETCONF requests (default 0.0.0.0).

Listen Port

Synopsis: *unsigned short integer*

Default: *830*

The port on which NETCONF listens for NETCONF requests. The default is port 830.

Extra IP:Ports

Synopsis: *A string*

Synopsis: *"extra-ip-ports" occurs in an array.*

Additional IP addresses and ports on which NETCONF listens for NETCONF requests. You can specify IP addresses and ports in the following forms:

- nnn.nnn.nnn.nnn:port represents an IPv4 address followed by a colon and port number. For example, 192.168.10.12:19343
- [::] represents the default IPv4 address and default port number. This is the default configuration.
- [::]:port represents an IPv6 address followed by a colon and port number. For example, [fe80::5eff:35ff]:16000

Maximum Number of NETCONF Sessions

Synopsis: *unsigned integer*

Synopsis: *- the keyword { unbounded }*

Default: *10*

The maximum number of concurrent NETCONF sessions

Idle Timeout

Default: *PT0S*

Maximum idle time before terminating a NETCONF session. If the session is waiting for notifications, or has a pending confirmed commit, the idle timeout is not used. The default value is 0, which means no timeout.

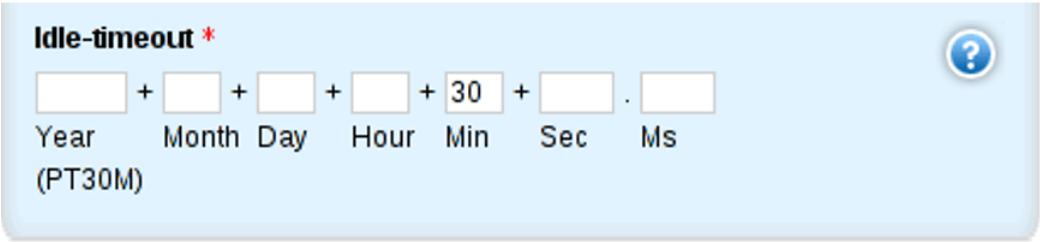

Figure 11.3. Idle-timeout field

Clicking on the Idle-timeout field on the NETCONF Sessions form allows you to choose a value for this field. The default value is PT30M, which stands for "Precision Time 30 Minutes". This refers to the time when an inactive session expires or times out. Only integer values corresponding to the following fields can be entered: Year, Month, Day, Hour, Min, Sec, or Ms. The example above shows the default value of PT30M, which corresponds to the Min field.

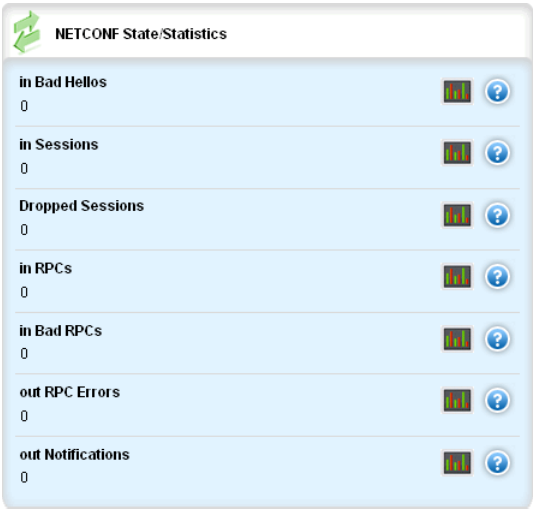

Figure 11.4. NETCONF State/Statistics form

in Bad Hellos

Synopsis: *unsigned integer*

The total number of sessions silently dropped because an invalid 'hello' message was received. This includes hello messages with a 'session-id' attribute, bad namespace, and bad capability declarations.

in Sessions

Synopsis: *unsigned integer*

The total number of NETCONF sessions started towards the NETCONF peer.

inSessions - inBadHellos = 'number of correctly started netconf sessions'

Dropped Sessions

Synopsis: *unsigned integer*

The total number of NETCONF sessions dropped.

inSessions - inBadHellos = 'number of correctly started netconf sessions'

in RPCs

Synopsis: *unsigned integer*

The total number of rpc requests received.

in Bad RPCs

Synopsis: *unsigned integer*

The total number of rpcs which were parsed correctly, but couldn't be serviced because they contained non-conformant XML.

out RPC Errors

Synopsis: *unsigned integer*

The total number of 'rpc-reply' messages with 'rpc-error' sent.

out Notifications

Synopsis: *unsigned integer*

The total number of 'notification' messages sent.

12. Chassis Management

This chapter describes basic system commands to control the ROX™-based system. Among the data available are:

- Hardware type and revision information
- Module firmware versions
- Current module status
- Temperature and power supply sensor data

This chapter also describes the “Chassis” menu and sub menus, which provide comprehensive diagnostic information for the RuggedRouter® chassis and all installed hardware modules.

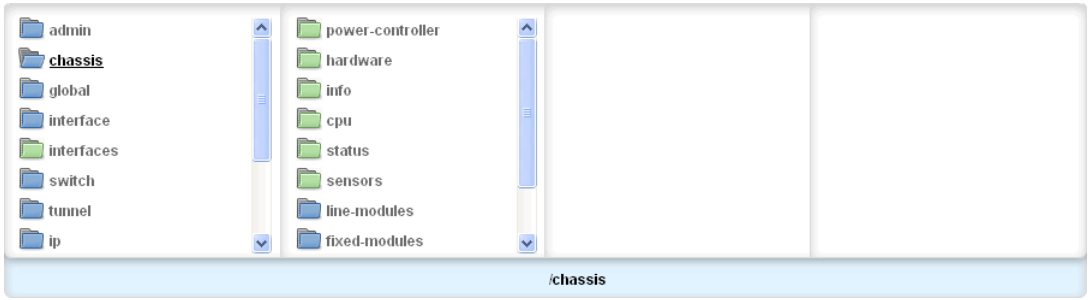

Figure 12.1. Chassis menu

The Chassis menu is accessible from the main menu under *Chassis*. This menu contains chassis-level configuration and status features. A variety of sub menus can be linked to from the Chassis menu. The Chassis sub menu section is organized so that information tables appear on a screen closer to the top level and clicking on one of the sub menus brings up the form(s) associated with the table. For example, clicking on the Chassis menu and then the Hardware sub menu will display the Slot Hardware table. Further clicking on the pm1 sub menu will display the Slot Hardware form.

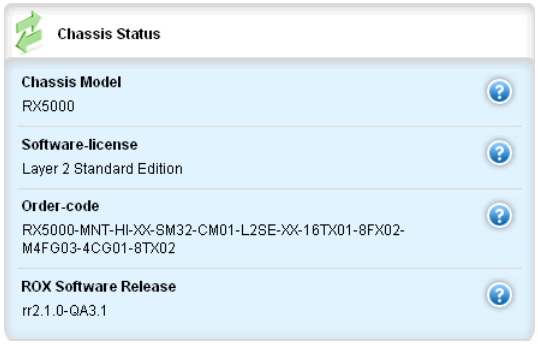

Figure 12.2. Chassis Status form

The Chassis Status form contains basic status information about the chassis. This form appears on the same screen as the *Chassis* menu.

Chassis Model

Synopsis: *string*

The RuggedCom device model name.

software-license

Synopsis: *string*

The current software capability.

order-code

Synopsis: *A string*

The order code derived from the current configuration of the device.

ROX Software Release

Synopsis: *string*

The release of ROX running on the chassis.

12.1. Power Controller

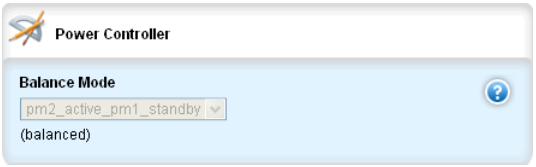

Figure 12.3. Power Controller form

As of ROX version 2.2, the balance-mode feature is not supported. This feature remains in the interface for backwards compatibility.

balance-mode

synopsis: string - one of the following keywords { PM2_Active_PM1_Standby, PM1_Active_PM2_Standby, Balanced }

When more than one power modules are present, this parameter specifies how they share the provision of power to the chassis. Select the "Balanced" option to cause each PM to provide an equal amount of power, Select the "PM1_Active_PM2_Standby" or "PM2_Active_PM1_Standby" options to provide the power from the active PM. The default value of this parameter is "Balanced"

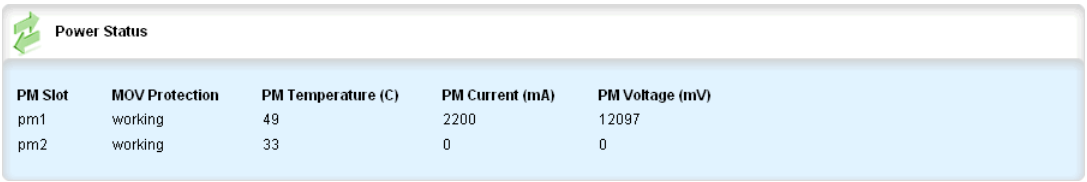

Figure 12.4. Power Status table

This table displays status information for power modules.

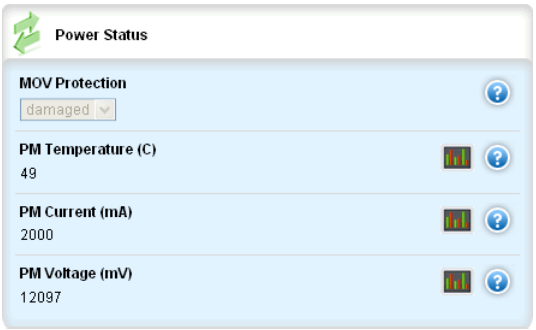

Figure 12.5. Power Status form

PM Slot

Synopsis: *string - one of the following keywords { pm2, pm1 }*

The name of the power module slot as labeled on the chassis

MOV Protection

Synopsis: *string - one of the following keywords { damaged, working, na }*

The state of the MOV protection circuit

PM Temperature (C)

Synopsis: *integer*

The temperature (Celsius) inside the power module

PM Current (mA)

Synopsis: *integer*

The current (mA) sourced by the power module

PM Voltage (mV)

Synopsis: *integer*

The voltage (mV) sourced by the power module

12.2. Slot Hardware

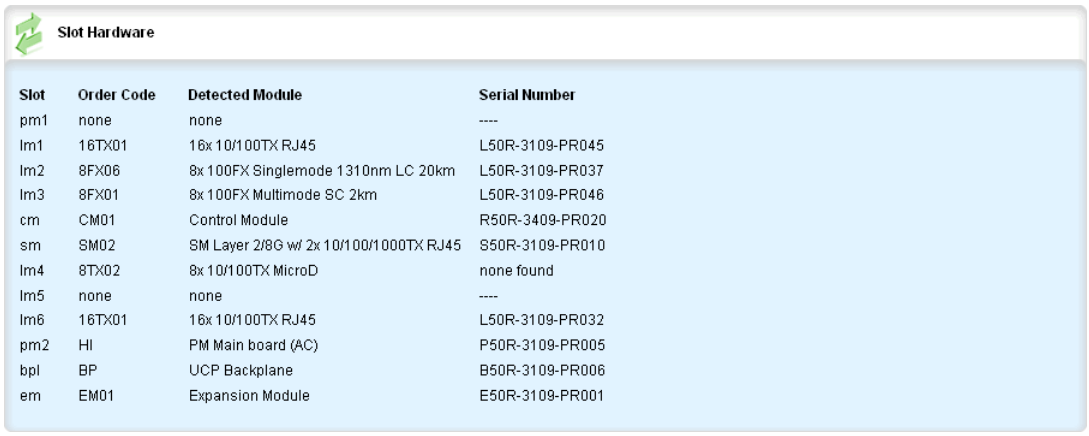
Slot Hardware
Slot Order Code Detected Module Serial Number
pm1 none none ----
lm1 16TX01 16x 10/100TX RJ45 L50R-3109-PR045
lm2 8FX06 8x 100FX Singlemode 1310nm LC 20km L50R-3109-PR037
lm3 8FX01 8x 100FX Multimode SC 2km L50R-3109-PR046
cm CM01 Control Module R50R-3409-PR020
sm SM02 SM Layer 2/8G w/ 2x 10/100/1000TX RJ45 S50R-3109-PR010
lm4 8TX02 8x 10/100TX MicroD none found
lm5 none none ----
lm6 16TX01 16x 10/100TX RJ45 L50R-3109-PR032
pm2 HI PM Main board (AC) P50R-3109-PR005
bpl BP UCP Backplane B50R-3109-PR006
em EM01 Expansion Module E50R-3109-PR001

Figure 12.6. Slot Hardware table

Figure 12.7. Slot Hardware form

The Slot Hardware table and form contain identifying information about the hardware module installed in a particular chassis slot.

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Order Code

Synopsis: *A string*

The order code of the chassis as derived from the current hardware configuration.

Detected Module

Synopsis: *A string*

The installed module's type specifier.

Serial Number

Synopsis: *A string*

The installed module's unique serial number.

12.3. Slot Identification

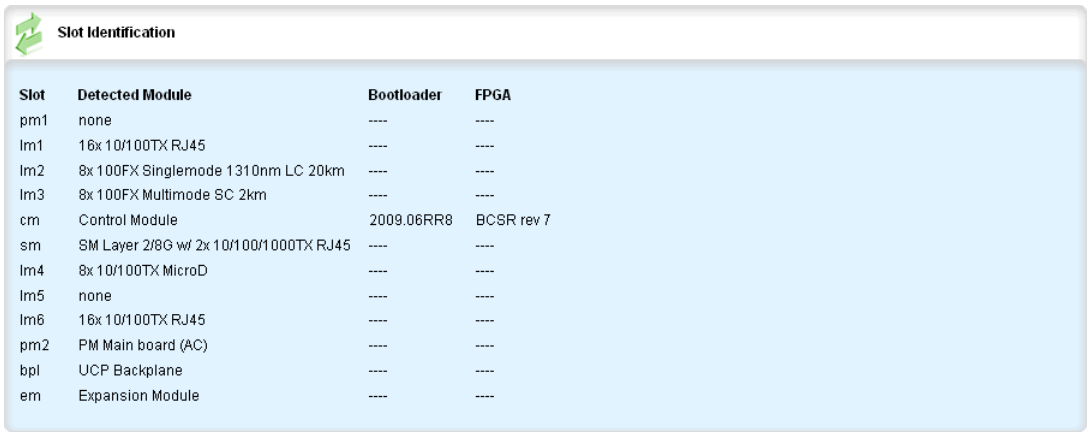
Slot Identification
Slot Detected Module Bootloader FPGA
pm1 none ---- ----
lm1 16x 10/100TX RJ45 ---- ----
lm2 8x 100FX Singlemode 1310nm LC 20km ---- ----
lm3 8x 100FX Multimode SC 2km ---- ----
cm Control Module 2009.06RR8 BCSR rev 7
sm SM Layer 2/8G w/ 2x 10/100/1000TX RJ45 ---- ----
lm4 8x 10/100TX MicroD ---- ----
lm5 none ---- ----
lm6 16x 10/100TX RJ45 ---- ----
pm2 PM Main board (AC) ---- ----
bpl UCP Backplane ---- ----
em Expansion Module ---- ----

Figure 12.8. Slot Identification table

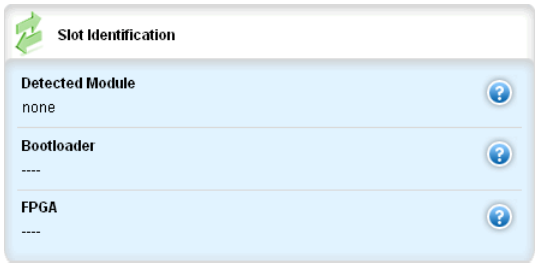

Figure 12.9. Slot Identification form

The Slot Identification table and form contain version information about the module software installed in a particular chassis slot.

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Detected Module

Synopsis: *A string*

The installed module's type specifier.

Bootloader

Synopsis: *string*

The version of the ROX bootloader software on the installed module.

FPGA

Synopsis: *string*

The version of the ROX FPGA firmware (if any) running on the installed module.

12.4. CPU

This section deals with CPU and memory usage for each slot (not all modules have a cpu). The Slot CPU table and form display status information about the hardware module installed in a particular chassis slot. The path to the Slot CPU/RAM Utilization table is *chassis/cpu*.

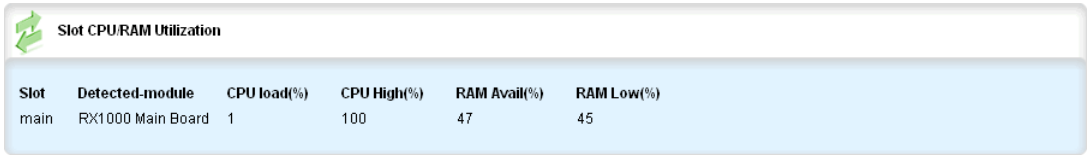

Figure 12.10. Slot CPU/RAM Utilization table

The path to the Slot CPU/RAM Utilization form is *chassis/cpu* and then clicking on a linked sub menu (for example, lm1).

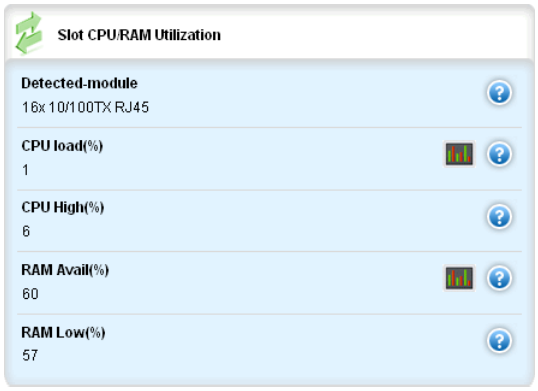

Figure 12.11. Slot CPU/RAM Utilization form

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

detected-module

Synopsis: *A string*

The installed module's type specifier.

CPU load(%)

Synopsis: *integer*

The CPU load, in percent, on the installed module.

RAM Avail(%)

Synopsis: *integer*

The proportion of memory (RAM) currently unused, in percent, on the installed module.

RAM Low(%)

Synopsis: *integer*

The lowest proportion of unused memory (RAM), in percent, recorded for the installed module since start-up.

12.5. Slot Status

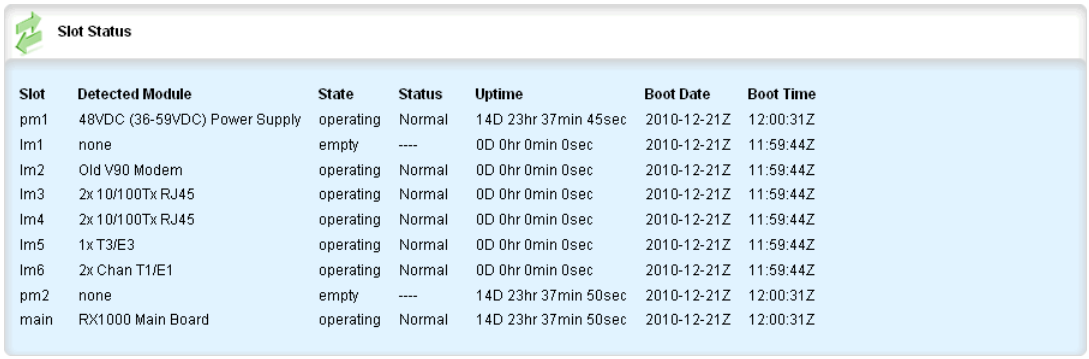
Slot Status
Slot Detected Module State Status Uptime Boot Date Boot Time
pm1 48VDC (36-59VDC) Power Supply operating Normal 14D 23hr 37min 45sec 2010-12-21Z 12:00:31Z
lm1 none empty ---- 0D 0hr 0min 0sec 2010-12-21Z 11:59:44Z
lm2 Old V90 Modem operating Normal 0D 0hr 0min 0sec 2010-12-21Z 11:59:44Z
lm3 2x 10/100Tx RJ45 operating Normal 0D 0hr 0min 0sec 2010-12-21Z 11:59:44Z
lm4 2x 10/100Tx RJ45 operating Normal 0D 0hr 0min 0sec 2010-12-21Z 11:59:44Z
lm5 1x T3/E3 operating Normal 0D 0hr 0min 0sec 2010-12-21Z 11:59:44Z
lm6 2x Chan T1/E1 operating Normal 0D 0hr 0min 0sec 2010-12-21Z 11:59:44Z
pm2 none empty ---- 14D 23hr 37min 50sec 2010-12-21Z 12:00:31Z
main RX1000 Main Board operating Normal 14D 23hr 37min 50sec 2010-12-21Z 12:00:31Z

Figure 12.12. Slot Status table

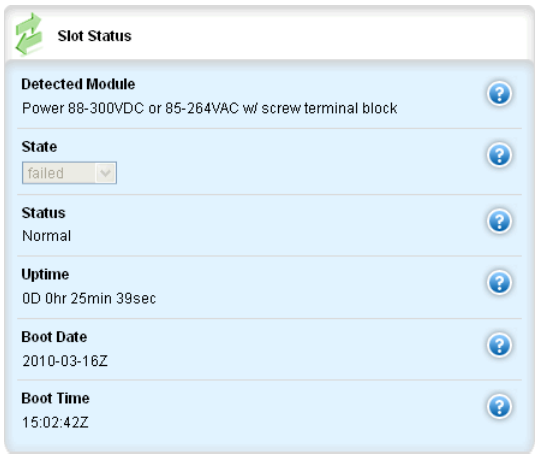

Figure 12.13. Slot Status form

The Slot Status table and form display status information about the hardware module installed in a particular chassis slot.

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Detected Module

Synopsis: *A string*

The installed module's type specifier.

State

Synopsis: *string - one of the following keywords { disconnected, failed, operating, resetting, disabled, empty, unknown }*

The current state of the installed module.

Status

Synopsis: *string*

The runtime status of the installed module.

Uptime

Synopsis: *string*

The total time elapsed since the start-up of the installed module.

Boot Date

Synopsis: *date specification*

The date on which the installed module was started up.

Boot Time

Synopsis: *time specification*

The time at which the installed module was started up.

12.6. Slot Sensors

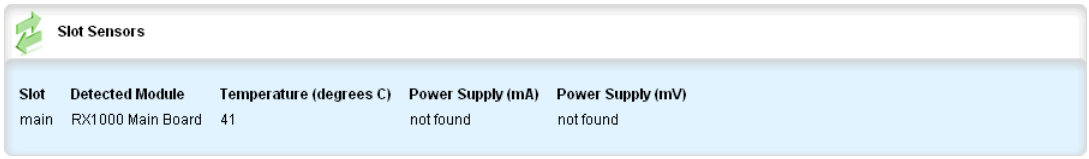

Figure 12.14. Slot Sensors table

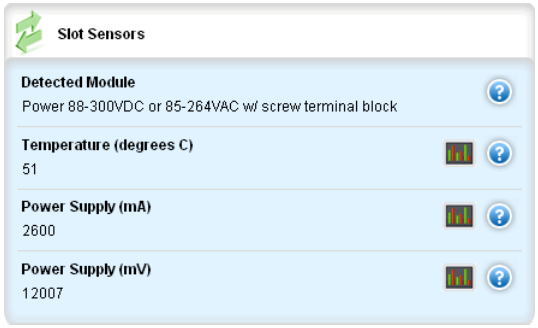

Figure 12.15. Slot Sensors form

Slot sensors contain temperature and power supply information about the installed modules.

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Detected Module

Synopsis: *A string*

The installed module's type specifier.

Temperature (degrees C)

Synopsis: *integer*

The temperature, in degrees C, of the installed module. If multiple temperature sensors are present on the board, the maximum reading is reported.

Power Supply (mA)

Synopsis: *integer*

The power supply current, in mA, being drawn by the installed module.

Power Supply (mV)

Synopsis: *integer*

The power supply voltage, in mV, seen by the installed module.

12.7. Module Configuration

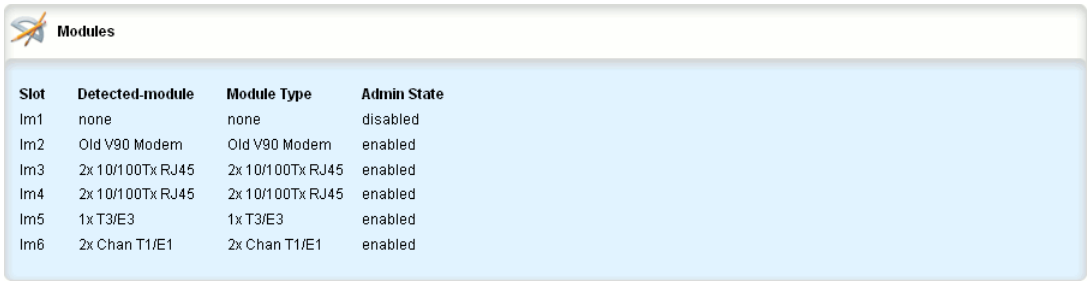
Modules
Slot Detected-module Module Type Admin State
lm1 none none disabled
lm2 Old V90 Modem Old V90 Modem enabled
lm3 2x 10/100Tx RJ45 2x 10/100Tx RJ45 enabled
lm4 2x 10/100Tx RJ45 2x 10/100Tx RJ45 enabled
lm5 1x T3/E3 1x T3/E3 enabled
lm6 2x Chan T1/E1 2x Chan T1/E1 enabled

Figure 12.16. Modules table

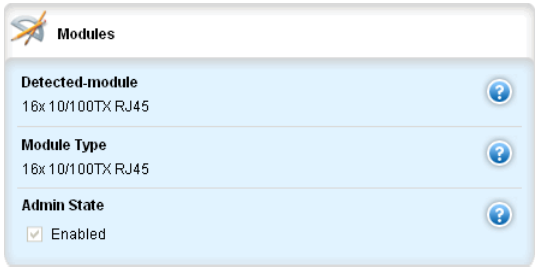

Figure 12.17. Modules form

The Module Configuration feature provides administrative control of the installed modules. The Modules table and form provide information about the administrative control of a module in a particular chassis slot.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The slot name, as marked on the silkscreen across the top of the chassis.

detected-module

Synopsis: *A string*

The installed module's type specifier.

Module Type

Synopsis: *A string*

Sets the module type to be used in this slot.

Admin State

Sets the administrative state for a module. Enabling the module powers it on.

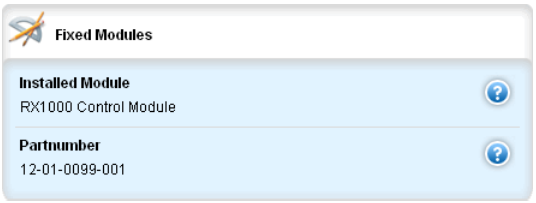

Figure 12.18. Fixed Modules form

Figure 12.19. Fixed Modules table

slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot name, as marked on the silkscreen across the top of the chassis.

Installed Module

Synopsis: *A string*

The module type to be used in this slot.

partnumber

Synopsis: *A string*

The part number of the module type in this slot.

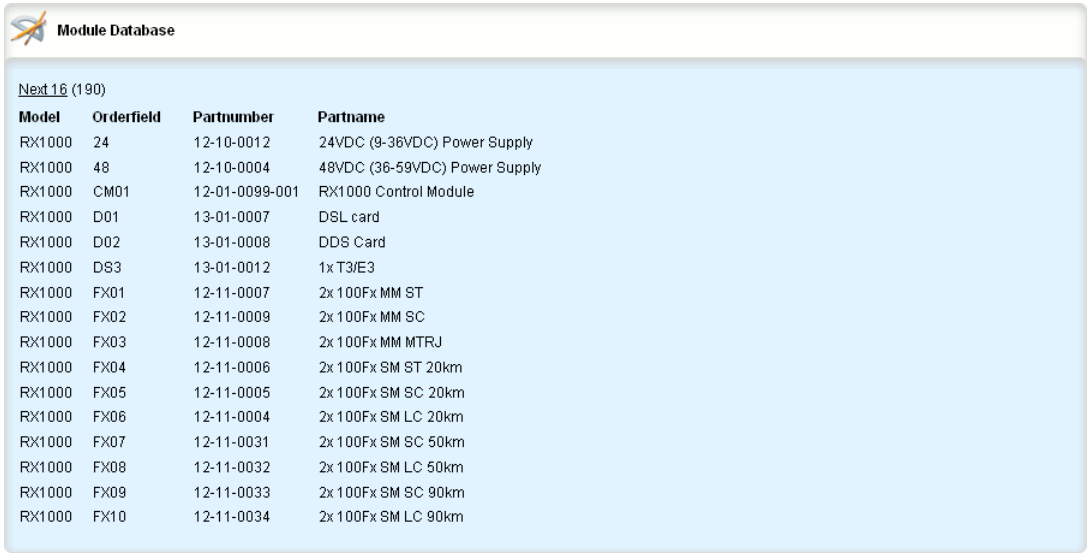

Module Database

Next 16 (190)

Model	Orderfield	Partnumber	Partname
RX1000	24	12-10-0012	24VDC (9-36VDC) Power Supply
RX1000	48	12-10-0004	48VDC (36-59VDC) Power Supply
RX1000	CM01	12-01-0099-001	RX1000 Control Module
RX1000	D01	13-01-0007	DSL card
RX1000	D02	13-01-0008	DDS Card
RX1000	DS3	13-01-0012	1x T3/E3
RX1000	FX01	12-11-0007	2x 100Fx MM ST
RX1000	FX02	12-11-0009	2x 100Fx MM SC
RX1000	FX03	12-11-0008	2x 100Fx MM MTRJ
RX1000	FX04	12-11-0006	2x 100Fx SM ST 20km
RX1000	FX05	12-11-0005	2x 100Fx SM SC 20km
RX1000	FX06	12-11-0004	2x 100Fx SM LC 20km
RX1000	FX07	12-11-0031	2x 100Fx SM SC 50km
RX1000	FX08	12-11-0032	2x 100Fx SM LC 50km
RX1000	FX09	12-11-0033	2x 100Fx SM SC 90km
RX1000	FX10	12-11-0034	2x 100Fx SM LC 90km

Figure 12.20. Module Database table

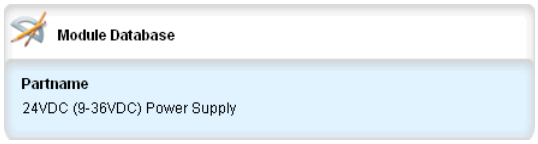

Module Database

Partname

24VDC (9-36VDC) Power Supply

Figure 12.21. Module Database form

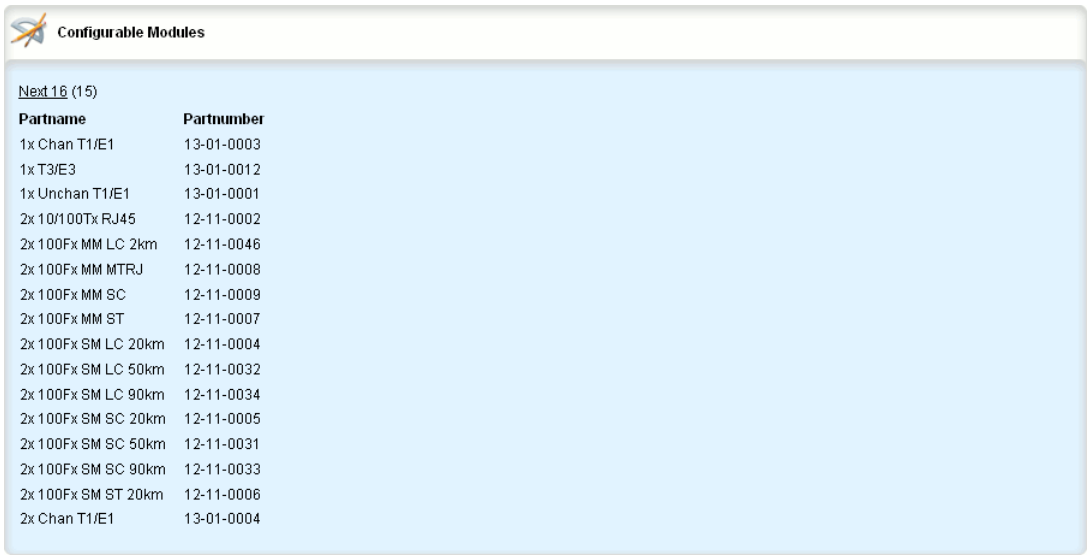

Configurable Modules

Next 16 (15)

Partname	Partnumber
1x Chan T1/E1	13-01-0003
1x T3/E3	13-01-0012
1x Unchan T1/E1	13-01-0001
2x 10/100Tx RJ45	12-11-0002
2x 100Fx MM LC 2km	12-11-0046
2x 100Fx MM MTRJ	12-11-0008
2x 100Fx MM SC	12-11-0009
2x 100Fx MM ST	12-11-0007
2x 100Fx SM LC 20km	12-11-0004
2x 100Fx SM LC 50km	12-11-0032
2x 100Fx SM LC 90km	12-11-0034
2x 100Fx SM SC 20km	12-11-0005
2x 100Fx SM SC 50km	12-11-0031
2x 100Fx SM SC 90km	12-11-0033
2x 100Fx SM ST 20km	12-11-0006
2x Chan T1/E1	13-01-0004

Figure 12.22. Configurable Modules table

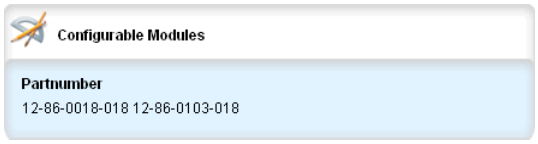

Configurable Modules

Partnumber

12-86-0018-018 12-86-0103-018

Figure 12.23. Configurable Modules form

13. PPP Users

13.1. Overview

Use the PPP menu to configure local and remote authentication for PPP user login through a dial-up modem and an L2TP client. An internal or external modem can be configured as a PPP Server or a PPP Client. For information on configuring a dial-up modem, see Chapter 18, *Modem*.

A PPP Server can be configured to accept a connection request only after validating the user's credentials. When so configured, the login credentials can be verified locally based on the information configured in *global/ppp/profiles/dial-in*, or can be verified through a remote RADIUS server configured in *global/ppp/radius*.

Use the dial-out option to create a profile for users who dial out to a PPP server. In this case, the device acts as a PPP client.

PPP users, profiles, and settings are configured on forms found under the PPP menu. To display the PPP menu, navigate to *global/ppp*.

13.2. PPP Configuration

The PPP menu is accessible from the main menu under *admin*. The PPP menu provides access to forms that allow you to add and remove PPP users and profiles and make adjustments to related settings.

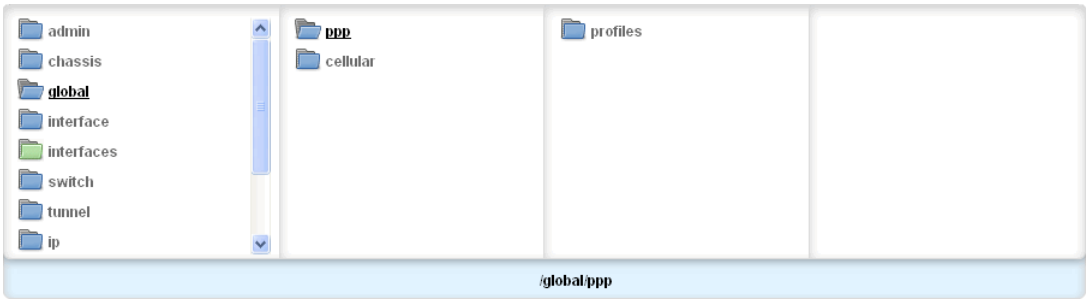

Figure 13.1. PPP menu

To display the Dial-in PPP Users table, navigate to *global/ppp/profiles/dialin*.

Figure 13.2. Dial-in PPP Users table

name

Synopsis: *A string*

The user ID of the remote client used to be authenticated

password

Synopsis: *A string*

The password of the remote client used to be authenticated

To display the Dial-in Users form, navigate to *global/ppp/profiles/dialin*, followed by any of the linked submenus.

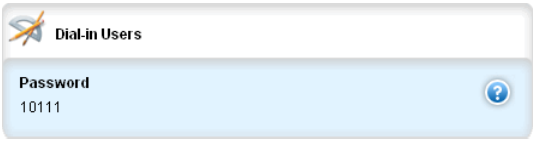

Figure 13.3. Dial-in Users form

The Dial-in Users form allows you to add PPP profiles for dial-in users.

To display the Dial-out PPP Users table, navigate to *global/ppp/profiles/dialout*.

Figure 13.4. Dial-out PPP Users table

Dial-out PPP is used to add PPP profile for dialOut users.

name

Synopsis: *A string*

The connection name

To display the PPP Configuration form, navigate to *global/ppp/profiles/dialout* and then click on any of the linked submenus.

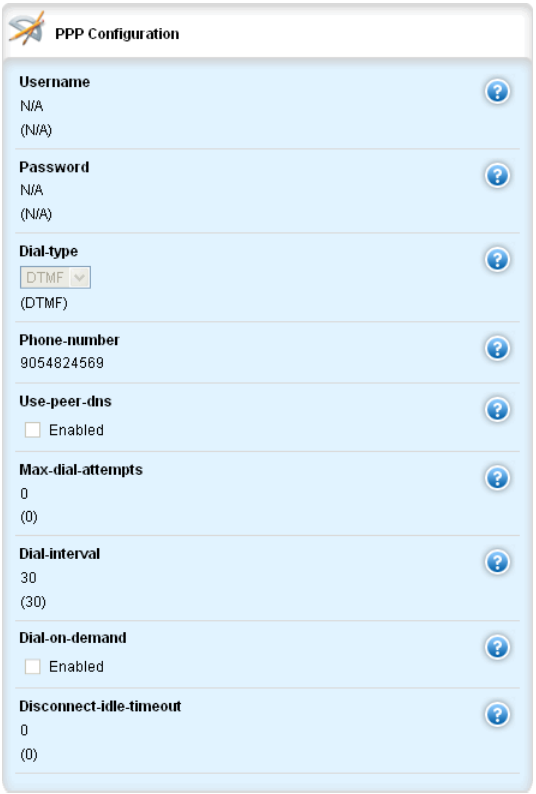

Figure 13.5. PPP Configuration form

username

Synopsis: *A string*

Default: *N/A*

The user ID used to log on to a remote PPP server

password

Synopsis: *A string*

Default: *N/A*

The password used to log on to a remote PPP server

dial-type

Synopsis: *string - one of the following keywords { Pulse, DTMF }*

Default: *DTMF*

The type of dialing system to use on the phone line. This field is only applied on analog modems

phone-number

Synopsis: *A string*

The telephone number to dial to connect to the remote PPP server

use-peer-dns

Using DNS recommended by the PPP server

max-dial-attempts

Synopsis: *integer*

Default:

The number of consecutive times that the modem will dial the phone number before it stops attempting to establish a connection. Zero (0) means the modem will try to connect to the PPP server forever.

dial-interval

Synopsis: *integer*

Default: *30*

The time in seconds to wait before re-initiating the link after it terminates

dial-on-demand

Activate Dial-on-Demand on this connection. The establishment of the PPP connection is postponed until there is data to be tranmitted via the interface

disconnect-idle-timeout

Synopsis: *integer*

Default:

The time in seconds to wait before disconnecting PPP when there is no traffic on the link. This option is only valid when Dial-on-Demand is enabled

failover-on-demand

Activate link failover on-demand on this device. PPP establishment on this device is controlled by link failover

To display the PPP RADIUS configuration forms, navigate to *global/ppp/radius*.

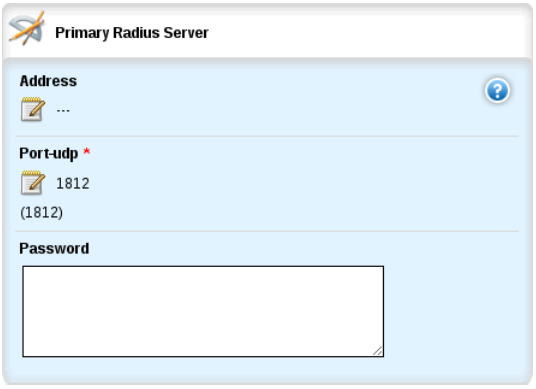

Figure 13.6. PPP Primary Radius Server form

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

port-udp

Synopsis: *integer*

Default: *1812*

password

Synopsis: *"AES CFB128"-encrypted string*

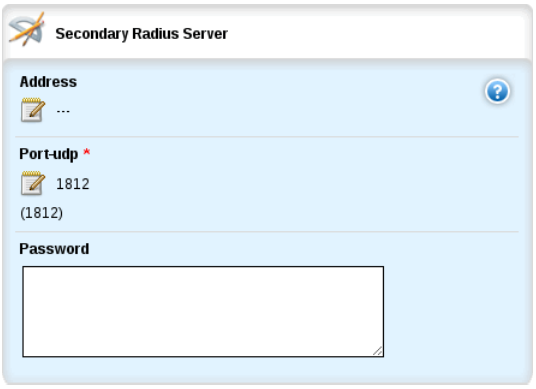

Figure 13.7. PPP Secondary Radius Server form

address

Synopsis: *IPv4 address in dotted-decimal notation*

The IPv4 address of the server

port-udp

Synopsis: *integer*

Default: *1812*

password

Synopsis: *"AES CFB128"-encrypted string*

13.3. PPP Interfaces and Link Failover

Link failover provides an easily configured means of raising a backup link upon the failure of a designated main link. A PPP interface can be specified as the backup link for a designate main link. For more information on link failover, see Chapter 29, *Link Failover*.

The PPP **Dial-on-demand** option is a standard PPP option. This option triggers the modem dial-out when there is traffic passing through the modem link. The modem hangs up when traffic stops within the time set in the PPP **Disconnect-idle-timeout** option. When **Dial-on-demand** is enabled, the presence of traffic controls the operation of the modem link.

When using a PPP interface with link failover, the link failover **On-demand** option allows link failover to bring up or take down the PPP interface as needed. Link failover triggers the modem dial-out to establish a PPP connection and pass traffic over the modem link when the main link is down. When the main link is up again, link failover takes down the PPP interface. The PPP **Disconnect-idle timeout** setting does not apply to the PPP interface when the interface is brought up by link failover.

*The link failover **On-demand** option and the PPP **Dial-on-demand** option are mutually exclusive: only one of these options should be set for a PPP interface.*

14. DHCP Server

14.1. DHCP Fundamentals

Dynamic Host Configuration Protocol (DHCP) is a method for centrally and consistently managing IP addresses and settings for clients, offering a variety of assignment methods. IP addresses can be assigned based on the Ethernet MAC address of a client, sequentially, or by using port identification provided by a DHCP relay agent device.

14.1.1. DHCP Network Organizations

The information to assign addresses in DHCP is organized to deal with clients at the interface, subnet, pool, shared network, host-group, and host levels.

Interfaces specify the IP interface to which the client sends a request.

Subnets control settings for each subnet that DHCP serves. A subnet can include a range of IP address to hand out to clients. Subnets contain groups, pools and hosts. Only one subnet can contain dynamic IP address ranges without any access restrictions on any given physical port, since DHCP doesn't know which subnet a client should belong to when the request is received.

Pools contain ranges of IP addresses to hand out to clients with access rules to determine which clients should receive addresses from that pool.

Shared networks are used when multiple subnets should be served by a single physical port. This applies both when using a DHCP relay agent connected to the port with additional subnets behind the relay agent, or when multiple virtual networks exist on one physical interface. Each subnet then gets its own subnet definition inside the shared network rather than at the top level. Shared networks contain subnets, groups and hosts.

Host-groups allow identical settings to be created for a group of hosts, making it easier to manage changes to the settings for all the hosts contained within the group. Groups contain hosts.

Host entries assign settings to a specific client based on its Ethernet MAC address.

14.1.2. Option 82 Support with Disable NAK

If DHCP relay clients (option 82 clients) are used on the same subnet as the DHCP server, some clients will try to renew a lease immediately after receiving it by requesting a renewal directly from the DHCP server. Because the DHCP server is only configured to provide the lease through a relay agent configured with the correct Option 82 fields, the server sends the client an NAK to disallow the lease. Enabling Option 82 disables the reject message so that the renewal request sent from the DHCP relay agent (which the DHCP server accepts since it has the correct Option 82 fields added) is the only message for which the client receives a reply. If the DHCP server and clients are not on the same subnet, this option is not required.

Note that the meaning of the value of many fields depends on the client's interpretation of the field, so the actual meaning of a field is determined by the client. To determine which values are required by the client for special options, refer to the client documentation.

14.2. Configuring DHCP Server

The DHCP Server menu is available under *services* at *services/dhcpserver*.

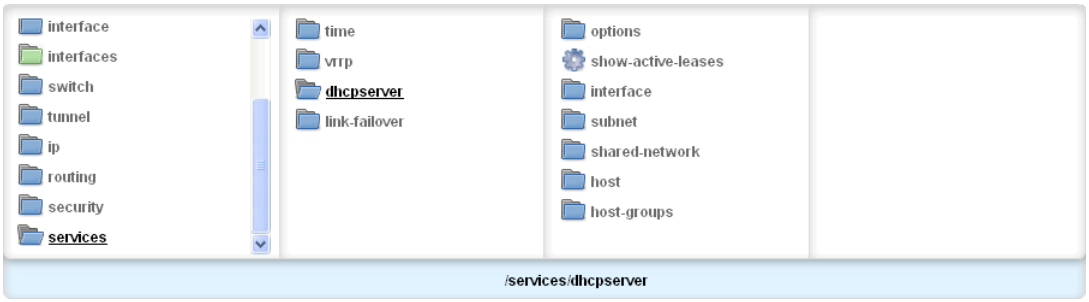

Figure 14.1. DHCP Server menu

Under *services/dhcpserver*, you can configure the following:

- enable the DHCP service. See Section 14.2.1, “Enabling the DHCP Service”.
- set the interface. See Section 14.2.2, “DHCP Interfaces”.
- configure subnets and pools. See Section 14.2.3, “DHCP Subnets and Pools”.
- configure shared networks. See Section 14.2.4, “DHCP Shared Networks”.
- configure hosts. See Section 14.2.5, “DHCP Hosts”.
- configure host-groups. See Section 14.2.6, “DHCP Host-groups”.
- configure DHCP options. See Section 14.2.8, “DHCP Options”.

Under *services/dhcpserver*, you can also view a list of active DHCP leases. See Section 14.2.7, “Viewing Active DHCP Leases”.

14.2.1. Enabling the DHCP Service

To enable or disable the DHCP service, navigate to *services/dhcpserver*. On the *Dynamic Host Control Protocol (DHCP) server* form, enable or disable the DHCP service.

Figure 14.2. DHCP Server form

enabled

Enables and disables the the DHCP server.

14.2.2. DHCP Interfaces

A *DHCP listen interface* is the interface on which DHCP listens for lease requests. You can configure multiple DHCP interfaces.

- To view a list of DHCP listen interfaces, navigate to *services/dhcpserver/interface*.

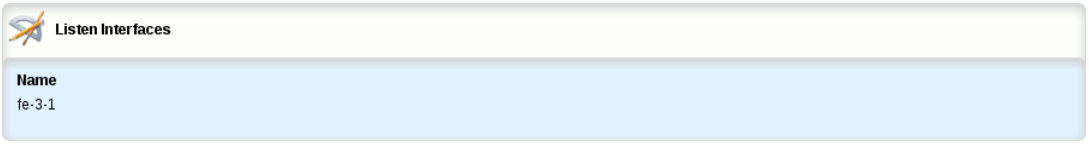

Figure 14.3. Listen Interfaces table

- To add a DHCP listen interface, enter edit mode, navigate to *services/dhcpserver/interface*, and click **<Add interface>**. On the *Key settings* form, select an interface from the list and click **Add**.

14.2.3. DHCP Subnets and Pools

- To view a list of DHCP subnets, navigate to */services/dhcpserver/subnet*.

Figure 14.4. Subnet Configuration table

- To add a subnet, enter edit mode, navigate to */services/dhcpserver/subnet*, and click **<Add subnet>**. On the *Key settings* form, type a name for the subnet and click **Add**.

  On the *Subnet Configuration* form, set the subnet parameters.

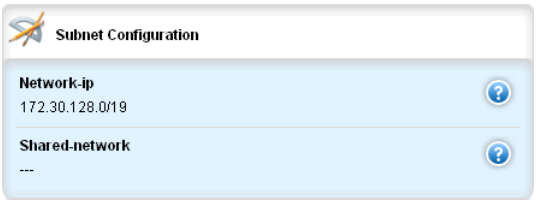

Figure 14.5. Subnet Configuration form

network-ip

Synopsis: *IPv4 address and prefix in CIDR notation*

The network IP address for this subnet

shared-network

Synopsis: *A string*

The shared-network that this subnet belongs to

You can configure DHCP options at the subnet level. Options set at this level override options set at higher levels.

- To set Lease Configuration and Client Configuration options, navigate to */services/dhcpserver/subnet{subnet id}/options*. For more information, see Section 14.2.8.1, “Lease Configuration Options” and Section 14.2.8.2, “Client Configuration Options at the DHCP Levels”.
- To set Client Configuration, NIS Configuration, and NetBios Configuration options, navigate to */services/dhcpserver/subnet{subnet id}/options/client*. For more information, see Section 14.2.8.3, “Client Configuration Options at the DHCP Client Level”.
- To set custom DHCP options, navigate to */services/dhcpserver/subnet{subnet id}/options/client/custom* and click <Add custom>. For more information, see Section 14.2.9, “Custom DHCP Options”.

14.2.3.1. DHCP Pools

- To view a list of DHCP pools, navigate to */services/dhcpserver/subnet{subnet02}/options/iprange*.

Figure 14.6. IP Pool Configuration table

- To add a DHCP pool, enter edit mode, navigate to */services/dhcpserver/subnet{subnet02}/options/iprange*, and click <Add iprange>. On the *Key settings* form, type the starting IP address of the range and click **Add**. On the *IP pool configuration* form, type the ending IP address of the range.

14.2.4. DHCP Shared Networks

- To view a list of shared networks, navigate to *services/dhcpserver/shared-network*.

Figure 14.7. Shared Network Configuration table

- To add a shared network, enter edit mode, navigate to *services/dhcpserver/shared-network*, and click **<Add shared-network>**. On the *Key settings* form, set a name for the shared network and click **Add**.

You can configure DHCP options at the shared network level. Options set at this level override options set at higher levels.

- To set Lease Configuration and Client Configuration options, navigate to */services/dhcpserver/shared-network{shared network id}/options*. For more information, see Section 14.2.8.1, “Lease Configuration Options” and Section 14.2.8.2, “Client Configuration Options at the DHCP Levels”.
- To set Client Configuration, NIS Configuration, and NetBios Configuration options, navigate to */services/dhcpserver/shared-network{shared network id}/options/client*. For more information, see Section 14.2.8.3, “Client Configuration Options at the DHCP Client Level”.
- To set custom DHCP options, navigate to */services/dhcpserver/shared-network{shared network id}/options/client/custom* and click <Add custom>. For more information, see Section 14.2.9, “Custom DHCP Options”.

14.2.5. DHCP Hosts

- To view a list of DHCP hosts, navigate to *services/dhcpserver/host*.

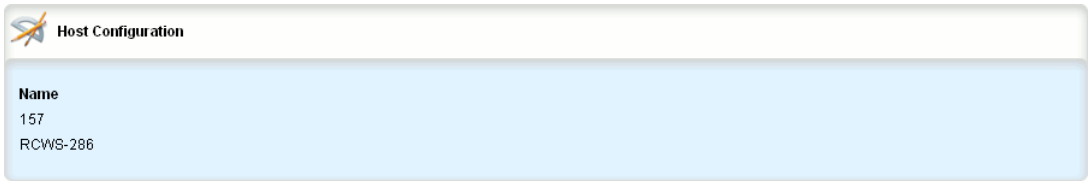

Figure 14.8. Host Configuration table

- To add a DHCP host, enter edit mode, navigate to *services/dhcpserver/host*, and click **<Add host>**. On the *Key settings* form, type a name for the host and click **Add**.

You can configure DHCP options at the host level. Options set at this level override options set at higher levels.

- To set Hardware Configuration, Lease Configuration, and Client Configuration options, navigate to */services/dhcpserver/host{host id}/options*. For more information, see Section 14.2.10, “Hardware Configuration”, Section 14.2.8.1, “Lease Configuration Options”, and Section 14.2.8.2, “Client Configuration Options at the DHCP Levels”.
- To set Client Configuration, NIS Configuration, and NetBios Configuration options, navigate to */services/dhcpserver/host{host id}/options/client*. For more information, see Section 14.2.8.3, “Client Configuration Options at the DHCP Client Level”.
- To set custom DHCP options, navigate to */services/dhcpserver/host{host id}/options/client/custom* and click <Add custom>. For more information, see Section 14.2.9, “Custom DHCP Options”.

14.2.6. DHCP Host-groups

- To view a list of host-groups, navigate to *services/dhcpserver/host-groups*.

Figure 14.9. Host Group Configuration table

- To add a host-group, enter edit mode, navigate to */services/dhcpserver/host-groups*, and click **<Add host-groups>**. On the *Key settings* form, type a name for the host-group and click **Add**.

You can configure DHCP options at the host-group level. Options set at this level override options set at higher levels.

- To set Lease Configuration and Client Configuration options, navigate to */services/dhcpserver/host-groups{host-group id}/options* . For more information, see Section 14.2.8.1, “Lease Configuration Options” and Section 14.2.8.2, “Client Configuration Options at the DHCP Levels”.
- To set Client Configuration, NIS Configuration, and NetBios Configuration options, navigate to */services/dhcpserver/host-groups{host-group id}/options/client*. For more information, see Section 14.2.8.3, “Client Configuration Options at the DHCP Client Level”.
- To set custom DHCP options, navigate to */services/dhcpserver/host-groups{host-group id}/options/client/custom* and click <Add custom>. For more information, see Section 14.2.9, “Custom DHCP Options”.

14.2.7. Viewing Active DHCP Leases

You can view a list of active DHCP leases. The list includes the start and end times, hardware ethernet address, and client hostname for each lease.

To view DHCP leases, navigate to *services/dhcpserver* and click **show-active-leases**. On the *Trigger Action* form, click **Perform**. The */services/dhcpserver/show-active-leases* form appears and displays the active DHCP leases.

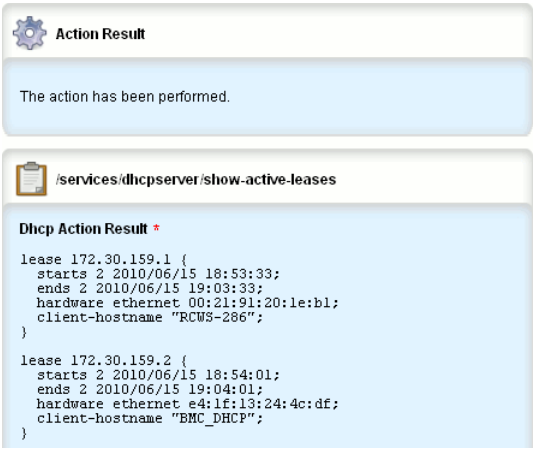

Figure 14.10. /services/dhcpserver/show-active-leases form

14.2.8. DHCP Options

You can set DHCP options at the subnet, shared network, host-groups, and hosts level. Options set at lower levels override those set at higher levels.

DHCP options are set on the following forms:

- *Leased Configuration* form: sets the DHCP lease times. This form is used at all DHCP levels.
- *Client Configuration* form at the DHCP grouping level: sets client address, host-group, and other settings. Subnets and shared networks use the same form; hosts and host-groups have unique settings on this form.
- *Client Configuration* form at the DHCP option level: sets hostname, subnet, DNS, and other settings. This form is used at all DHCP levels.
- *NIS Configuration* form: sets NIS server information. This form is used at all DHCP levels.
- *NetBios Configuration* form: sets NetBios scope and nameserver information. This form is used at all DHCP levels.
- *Custom* form: sets DHCP custom options. This form is used at all DHCP levels.
- *Hardware Configuration*: sets network type and MAC address information. This form is used for hosts only.

14.2.8.1. Lease Configuration Options

You can set the lease configuration options at all DHCP levels. To set DHCP lease configuration options, enter edit mode and navigate to:

- DHCP server options: */services/dhcpserver/options*
- subnet options: */services/dhcpserver/subnet{subnet id}/options*
- shared network options: */services/dhcpserver/shared-network{shared network id}/options*
- host-group options: */services/dhcpserver/host-groups{host-group id}/options*
- host options: */services/dhcpserver/host{host id}/options*

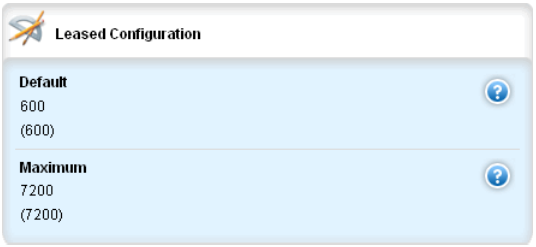

Figure 14.11. Lease Configuration form

default

Synopsis: *integer*

Default: *600*

The minimum leased time that the server offers to the client

maximum

Synopsis: *integer*

Default: *7200*

The maximum leased time that the server offers to the client

14.2.8.2. Client Configuration Options at the DHCP Levels

Different DHCP options are set at the subnet and shared network, hosts, and host groups levels.

14.2.8.2.1. Client Configuration Options: Subnets and Shared Networks

To set DHCP client configuration options at the subnet and shared networks levels, enter edit mode and navigate to:

- subnet options: */services/dhcpserver/subnet{subnet id}/options*
- shared network options: */services/dhcpserver/shared-network{shared network id}/options*

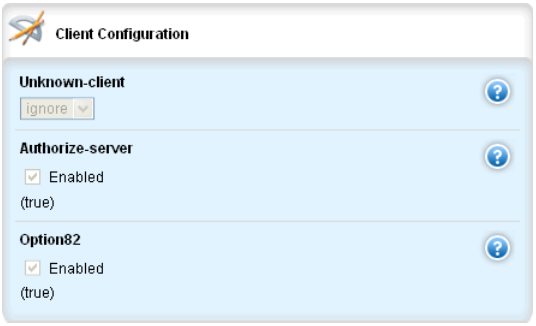

Figure 14.12. Client Configuration form for Subnets and Shared Networks

unknown-client

Synopsis: *string - one of the following keywords { ignore, deny, allow }*

The action to take for previously unregistered clients

authorize-server

Enables/disables the server's authorization on this client. If enabled, the server will send deny messages to the client that is trying to renew the lease, which the server knows the client shouldn't have

option82

Enable/disable the NAK of option 82 clients for this subnet

14.2.8.2.2. Client Configuration Options: Hosts

To set DHCP client configuration options at the host level, enter edit mode and navigate to */services/dhcpserver/host{host id}/options*.

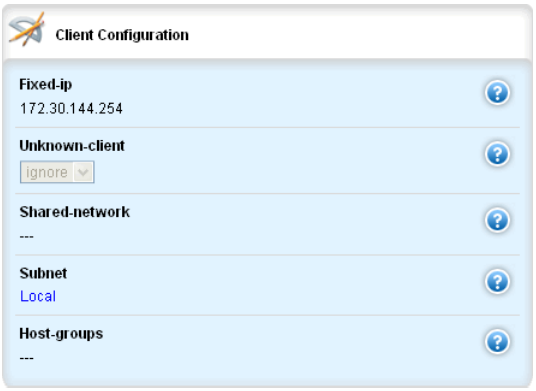

Figure 14.13. Client Configuration form for Hosts

fixed-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address that the server assigns to the matching client

unknown-client

Synopsis: *string - one of the following keywords { ignore, deny, allow }*

The action to take for previously unregistered clients

shared-network

Synopsis: *A string*

Shared-network that this host belongs to

subnet

Synopsis: *A string*

Subnet that this host belongs to

host-groups

Synopsis: *A string*

Host groups that this host belongs to

14.2.8.2.3. Client Configuration Options: Host-groups

To set DHCP client configuration options at the host-group level, enter edit mode and navigate to */services/dhcpserver/host-group{host-group id}/options*.

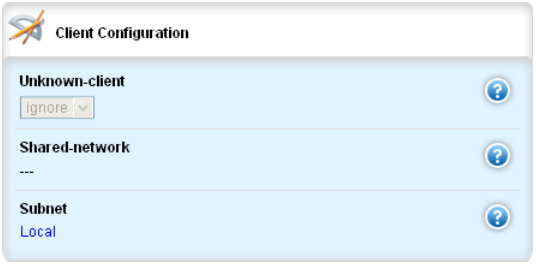

Figure 14.14. Client Configuration form for Host-groups

unknown-client

Synopsis: *string - one of the following keywords { ignore, deny, allow }*

Default: *allow*

The action to take for previously unregistered clients

shared-network

Synopsis: *A string*

Shared-network that this host group belongs to

subnet

Synopsis: *A string*

The subnet that this host group belongs to

14.2.8.3. Client Configuration Options at the DHCP Client Level

The following DHCP client, NIS, and NetBios options are set at the client level under all DHCP levels. To set the client configuration options, enter edit mode and navigate to:

- DHCP server options: */services/dhcpserver/options/client*
- subnet options: */services/dhcpserver/subnet{subnet id}/options/client*
- shared network options: */services/dhcpserver/shared-network{shared network id}/options/client*
- host-group options: */services/dhcpserver/host-groups{host-group id}/options/client*
- host options: */services/dhcpserver/host{host id}/options/client*

Client Configuration:

Figure 14.15. Client Configuration form for DHCP Clients

hostname

Synopsis: *A string*

The unique name to refer to the host within a DHCP configuration

subnetmask

Synopsis: *IPv4 address in dotted-decimal notation*

Subnet mask

default-route

Synopsis: *IPv4 address in dotted-decimal notation*

The default route that the server offers to the client when it issues the lease to the client

broadcast

Synopsis: *IPv4 address in dotted-decimal notation*

The broadcast address that the server offers to the client when it issues the lease to the client

domain

Synopsis: *string*

The domain name that the server offers to the client when it issues the lease to the client

dns-server

Synopsis: *IPv4 address in dotted-decimal notation*

The domain name server that the server offers to client when it issues the lease to client

static-route

Synopsis: *IPv4 address in dotted-decimal notation*

The static route that the dhcpserver offers to the client when it issues the lease to the client

NIS Configuration:

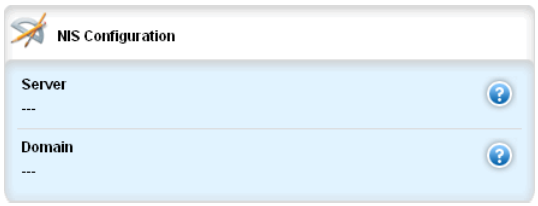

Figure 14.16. NIS Configuration form

server

Synopsis: *IPv4 address in dotted-decimal notation*

The NIS server address that the dhcpserver offers to the client when it issues the lease to the client

domain

Synopsis: *string*

The NIS domain name that the dhcpserver offers to the client when it issues the lease to the client

NetBios Configuration:

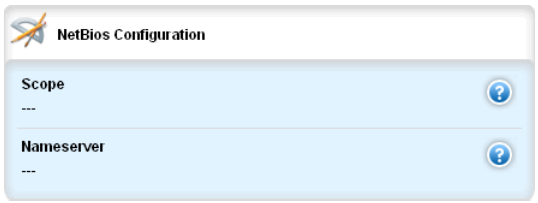

Figure 14.17. Netbios Configuration form

This is the NetBIOS nameserver that dhcpserver offers to client when it issues the lease to a client.

scope

Synopsis: *string*

Default: *netbios*

The NetBIOS scope that the dhcpserver offers to the client when it issues the lease to the client

nameserver

Synopsis: *string*

Default: *127.0.0.1*

The NetBIOS nameserver that the dhcpserver offers to the client when it issues the lease to the client

14.2.9. Custom DHCP Options

You can set custom DHCP options at the under clients at all DHCP levels. To set a custom DHCP option, you need to know the number of the option you want to set and the valid values for the option.

To set a custom DHCP option, enter edit mode and navigate to one of the following locations and click **<Add custom>**:

- DHCP server options: */services/dhcpserver/options/client/custom*
- subnet options: */services/dhcpserver/subnet{subnet id}/options/client/custom*
- shared network options: */services/dhcpserver/shared-network{shared network id}/options/client/custom*
- host-group options: */services/dhcpserver/host-groups{host-group id}/options/client/custom*
- host options: */services/dhcpserver/host{host id}/options/client/custom*

On the Key settings form, set the number and value for the DHCP custom option.

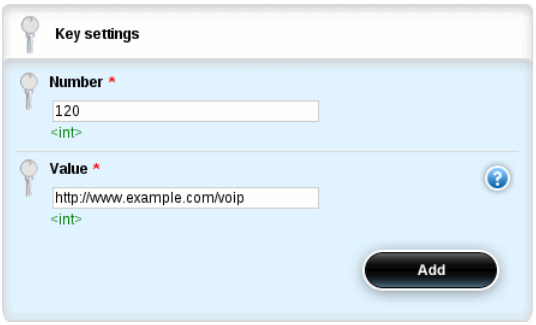

Figure 14.18. Setting a DHCP Custom Option

14.2.10. Hardware Configuration

The *Hardware Configuration* form is only available for DHCP hosts. To set the hardware options for a DHCP host, enter edit mode and navigate to */services/dhcpserver/host{host_01}/options*.

Figure 14.19. Hardware Configuration form

type

Synopsis: *string - one of the following keywords { ethernet, token-ring, fddi }*

Default: *ethernet*

The type of network hardware used by the client, associated with the host entry.

mac

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The physical network address of the client. Note that this corresponds to the hardware type; for example, MAC address for ethernet.

Part II. Network Interfaces and Ethernet Bridging

This part describes network interfaces and the configuration and monitoring of Ethernet bridging on a ROX™-based networking device:

Ethernet Statistics	Chapter 15, *Ethernet Statistics*
IP Statistics	Chapter 16, *IP Statistics*
Virtualswitch Bridging	Chapter 17, *Virtual Switch Bridging*
Modem	Chapter 18, *Modem*
Serial Ports	Chapter 19, *Serial Protocols*
WAN	Chapter 20, *WAN*

15. Ethernet Statistics

ROX™ provides the following features for gathering and reporting Ethernet statistics:

- Viewing basic Ethernet statistics.

15.1. Viewing Non-switched Ethernet Statistics

The Non-switched Ethernet Statistics menus are accessible from the main menu under *Interfaces*. Statistics can be found under the submenus that follow. For instance, click on *eth* and then click on a linked submenu for an interface (for example, fe-3-1).

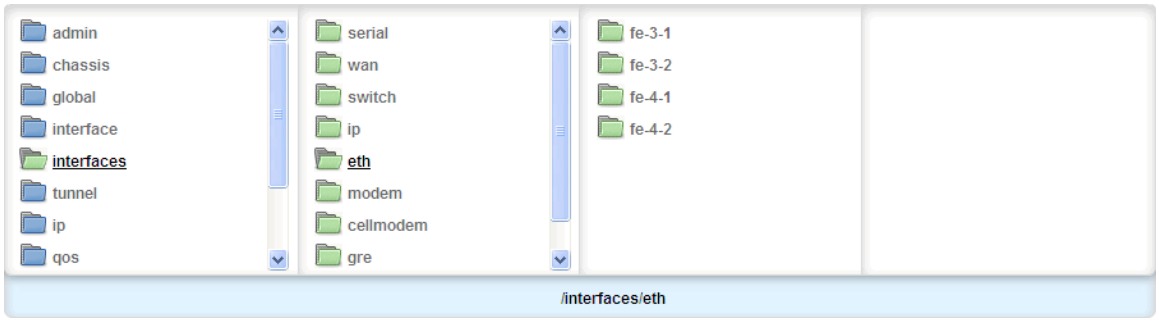

Figure 15.1. Statistics Menu

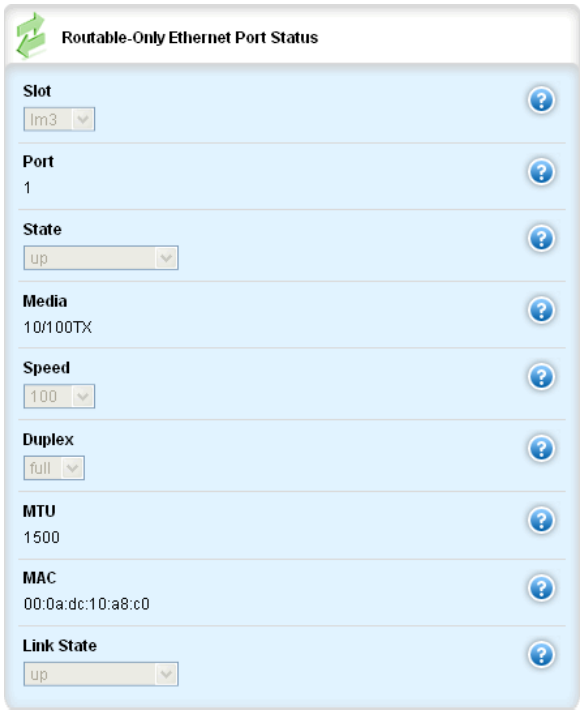

Figure 15.2. Routable-Only Ethernet Port Status Form

The Routable-Only Ethernet Port Status, Receive Statistics, and Transmit Statistics forms appear on the same screen as the *Statistics* menus. The Routable Ethernet Ports form displays the ethernet port configuration and status for a port. Ethernet statistics for the system's IP interfaces are available on the Receive Statistics and Transmit Statistics forms.

Name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

Slot

Synopsis: *string - the keyword { --- }*

Synopsis: *string - one of the following keywords { main, pm2, pm1 }*

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string - one of the following keywords { em, cm }*

Synopsis: *string - the keyword { trnk }*

The slot of the module that contains this port.

Port

Synopsis: *integer*

The port number on the module.

state

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

The port's link status.

media

Synopsis: *A string*

The type of port media { 100TX, 10FL, 100FX, 1000X, 1000T, 802.11g, EoVDSL, 100TX }. It provides the user with a description of the installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), and may be Short Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, the media description is displayed per the SFF-8472 specification, if the transceiver is plugged into the module. E.g. 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST,1000SX SFP LC S SL M5.

Speed

Synopsis: *string - one of the following keywords { 1000, 100, 10 }*

Link speed (in Megabits-per-second or Gigabits-per-second)

Duplex

Synopsis: *string - one of the following keywords { full, half }*

Link duplex status.

MTU

Synopsis: *integer*

MTU (Maximum Transmission Unit) value on the port.

MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The MAC address of the port.

Link State

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

Link status.

Figure 15.3. Receive Statistics Form

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the receive device.

Figure 15.4. Transmit Statistics Form

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets transmitted.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the transmit device.

Collisions

Synopsis: *unsigned integer*

Number of collisions detected on the port.

16. IP Statistics

The forms and tables accessible from the Interfaces IP menu (below) show the status of what has been configured using the forms and tables from the Interface and IP menus.

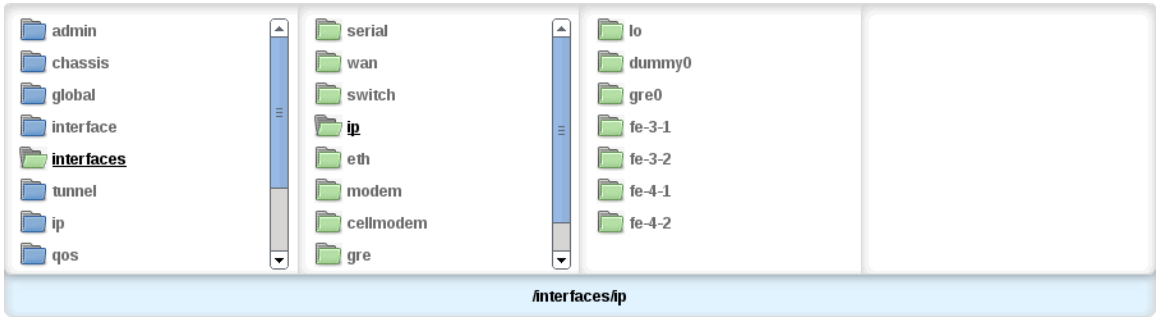

Figure 16.1. Interfaces IP Menu

The Interfaces IP menu is accessible from the main menu under *interfaces/ip*.

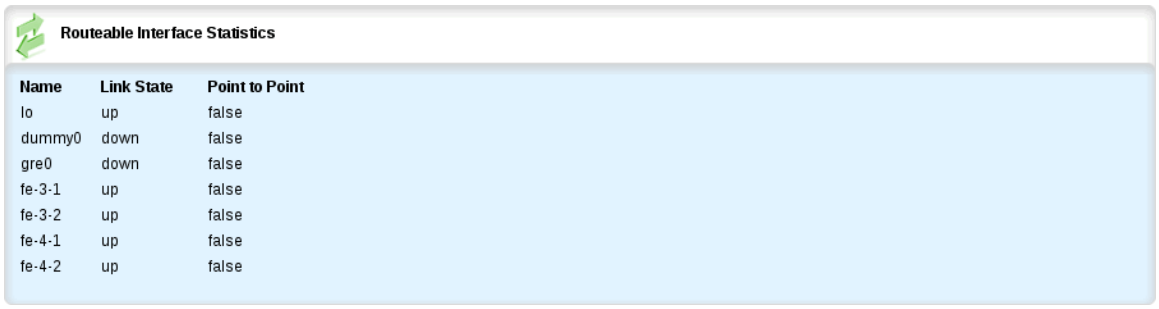

Figure 16.2. Routable Interface Statistics Table

This table appears on the same screen as the *Interfaces IP* menu.

The path to the Routable Interface Statistics form, Receive Statistics form and Transmit Statistics form is *interfaces/ip/{interface}*.

Figure 16.3. Routable Interface Statistics Form

Name

Synopsis: *A string*

The interface name

Link State

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

Up or Down

Point to Point

Synopsis: *boolean*

Is point to point link.

Figure 16.4. Receive Statistics Form

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the receive device.

Figure 16.5. Transmit Statistics Form

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets transmitted.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the transmit device.

Collisions

Synopsis: *unsigned integer*

Number of collisions detected on the port.

17. Virtual Switch Bridging

17.1. Overview

A virtual switch bridges different network segments in way that is not dependent on a particular protocol.

Network traffic between segments is forwarded regardless of the IP and MAC addresses in a packet. In a virtual switch, forwarding is done in Layer 2 and allows all network traffic, including L2 Multicast (GOOSE, ISO), IP Multicast, Unicast, and Broadcast messages, to go through the virtual switch tunnel without any modifications. A virtual switch can be useful for GOOSE messaging when the sender and receiver need to communicate through a routable IP network. Because there is no IP encapsulation for the L2 traffic going through the virtual switch, network latency is minimized for the traffic between end devices.

The virtual switch appears on the device as a virtual Ethernet interface over a physical interface (Ethernet port) between two routers. Physically, the two routers can be in different locations. There can be multiple virtual switch instances in a router. Each instance can include two or more interfaces, but an interface can only be a member of one virtual switch instance.

A virtual switch interface in a router can be a routable interface when an IP address is assigned either statically or via DHCP. The network address assigned to the virtual switch interface can be included in the dynamic routing protocol and the interface can carry a routing update. The IP address assigned to the virtual switch can be used as the default gateway for the end devices connected to the virtual switch interface. Network services, such as SSH, DHCP, NTP, VRRP, etc, can be configured to run on the virtual switch interface.

17.1.1. Helpful Hints

- Any IP address assigned to an interface becomes inactive and hidden when the interface is added to the virtual switch. The address on the interface is reactivated after removing the interface from the virtual switch.
- Be careful when adding interfaces to the virtual switch. Any network services running on the individual interfaces will need to be reconfigured after adding the interface to the virtual switch. For example, if a DHCP server running on fe-3-1 and fe-3-1 is subsequently made a member of the VirtualSwitch VS1, the DHCP configuration must be changed to refer to VS1.
- In ROX™, the virtual switch is implemented in the software. Therefore, a CPU resource is needed to perform forwarding of broadcast, multicast and unicast traffic.
- If the router is running as a firewall, the routeback option must be enabled for the virtual switch interface in the “fwinterface” submenu under the Firewall menu.

17.2. Sample Use Case

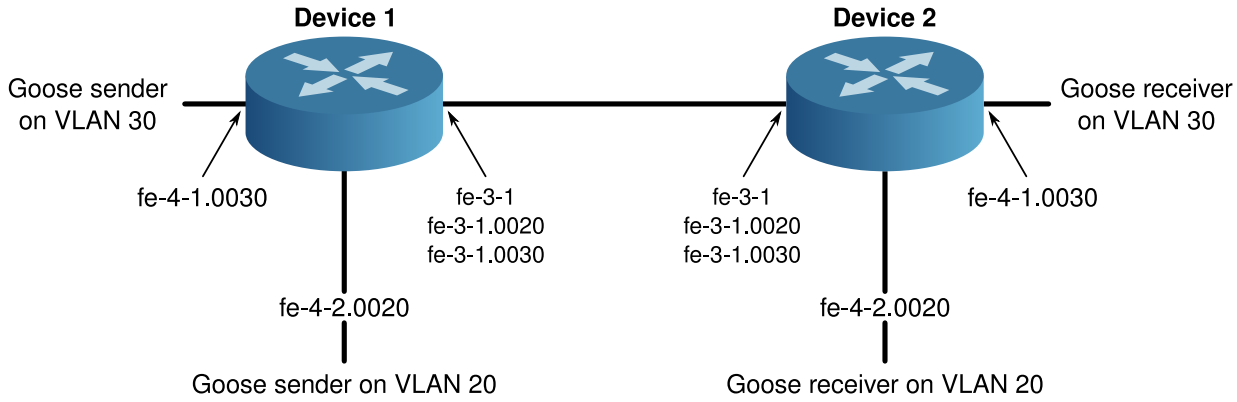

Figure 17.1. Virtual switch with multiple interfaces

To create the configuration shown in this example, follow these steps:

1. Configure the port connected to the senders and receivers as follows:
   - PVID 20.
   - PVID 30.
2. Configure the Ethernet interface, fe-3-1, on both devices with two VLANs: VLAN 20 and VLAN 30.
3. Configure two instances of VirtualSwitch by adding the following interfaces to the virtual switch on both devices:
   - VS1 on Device 1: fe-4-2.0020, fe-3-1.0020
   - VS2 on Device 1: fe-4-1.0300, fe-3-1.0030
4. Use the same configuration for Device 2.
5. Assign IP addresses to the virtual switch instances on both the devices:
   - VS1 on Device 1: 192.168.11.11/24
   - VS2 on Device 1: 192.168.22.22/24
   - VS1 on Device 2: 192.168.11.12/24
   - VS2 on Device 2: 192.168.22.23/24

When configuration is complete, tagged or untagged traffic received on VS1 of Device 1 should only be forwarded to VS1 on Device 2. Similarly, traffic received on VS2 of Device 1 should only be forwarded to VS2 on Device 2.

17.3. Virtual Switch Configuration and Status

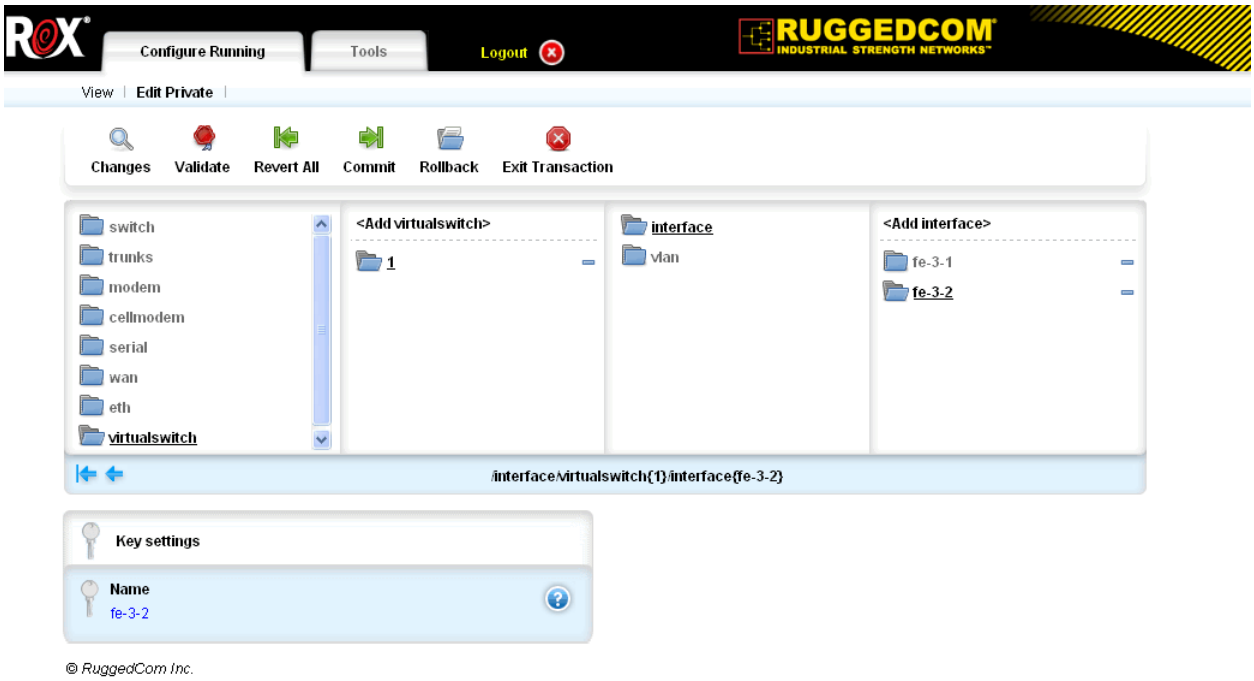

Figure 17.2. Adding a Virtual Switch

To add a virtual switch, enter Edit Private mode. Add a virtual switch and at least two interfaces. You can also add VLANs.

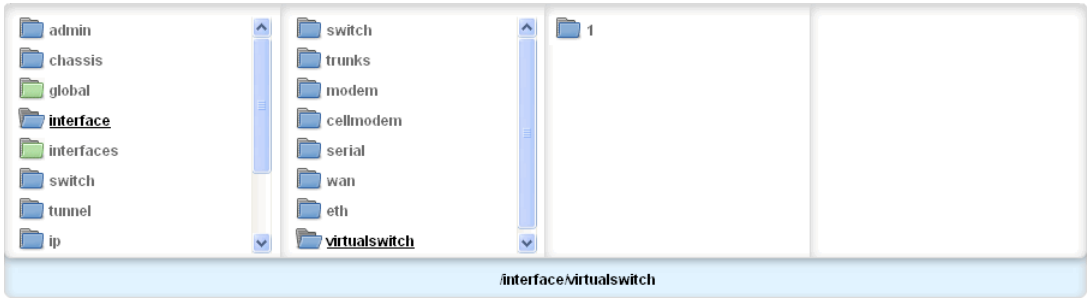

Figure 17.3. Interface Virtualswitch menu

The Interface Virtualswitch menu is located at *interface/virtualswitch*. If a virtual switch is configured, the Virtualswitch table appears on the same screen as this menu.

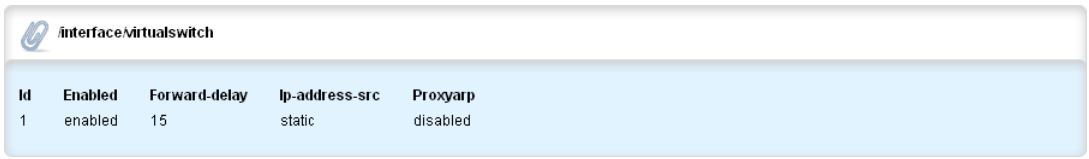

Figure 17.4. Virtualswitch table

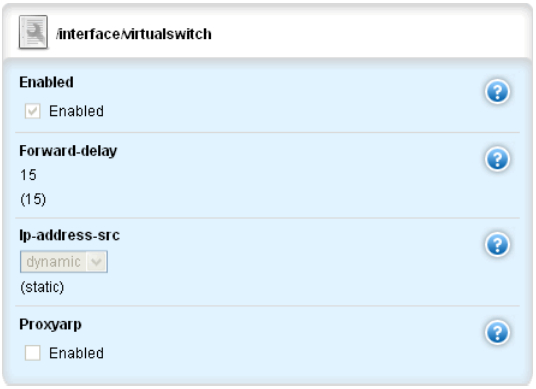

Figure 17.5. Virtualswitch form

To display this form, navigate to *interface/virtualswitch/{number}*.

Forward Delay

Synopsis: *unsigned byte*

Default: *15*

Delay (in seconds) of the listening and learning state before goes to forwarding state.

Alias

Synopsis: *A string*

The SNMP alias name of the interface

IP Address Source

Synopsis: *string - one of the following keywords { dynamic, static }*

Default: *static*

Whether the IP address is static or dynamically assigned via DHCP or BOOTP.

ProxyARP

Enables/Disables whether the port will respond to ARP requests for hosts other than itself

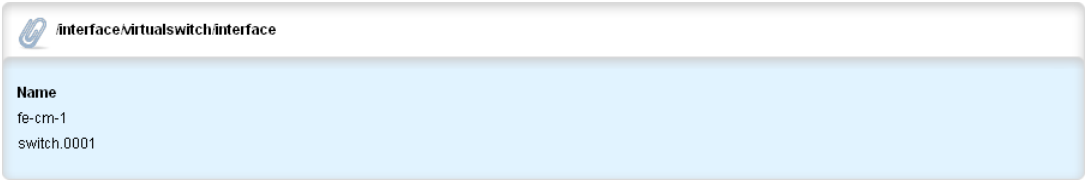

Figure 17.6. Interface table

To display this table, navigate to *interface/virtualswitch/{number}/interface*.

Interface Name

Synopsis: *A string*

Interface name.

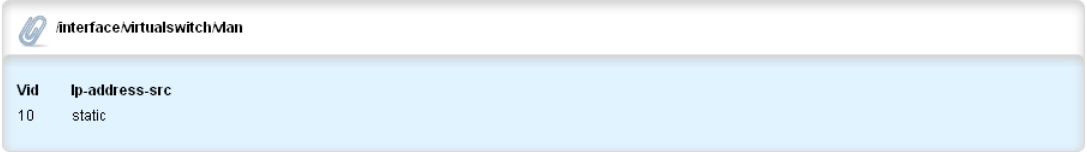

Figure 17.7. VLAN table

To display this table, navigate to *interface/virtualswitch/{number}/vlan*.

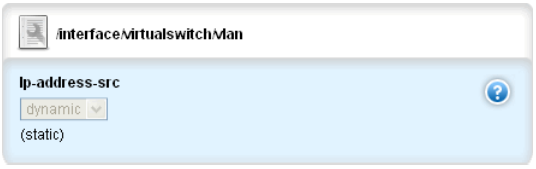

Figure 17.8. VLAN form

To display this form, navigate to *interface/virtualswitch/{number}/vlan/{number}*.

VLAN ID

Synopsis: *integer*

VLAN ID for this routable logical interface

IP Address Source

Synopsis: *string - one of the following keywords { dynamic, static }*

Default: *static*

Whether the IP address is static or dynamically assigned via DHCP or BOOTP.

If a virtual switch has been configured, some virtual switch data will be displayed under the Interfaces Virtualswitch menu.

Figure 17.9. Interfaces Virtualswitch menu

To display the Interfaces Virtualswitch menu, navigate to *interfaces/virtualswitch*. The Virtualswitch table appears on the same screen as this menu.

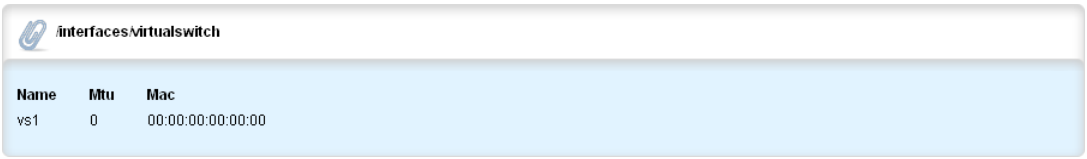

Figure 17.10. Virtualswitch table

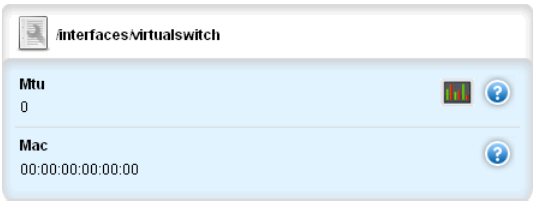

Figure 17.11. Virtualswitch form

To display the Virtualswitch, Receive, and Transmit forms, navigate to *interfaces/virtualswitch/{virtualswitch number}*.

Interface Name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

MTU

Synopsis: *integer*

MTU (Maximum Transmission Unit) value on the port.

MAC

Synopsis: *Ethernet MAC address in colon-separated hexadecimal notation*

The MAC address of the port.

Figure 17.12. Receive form

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the receive device.

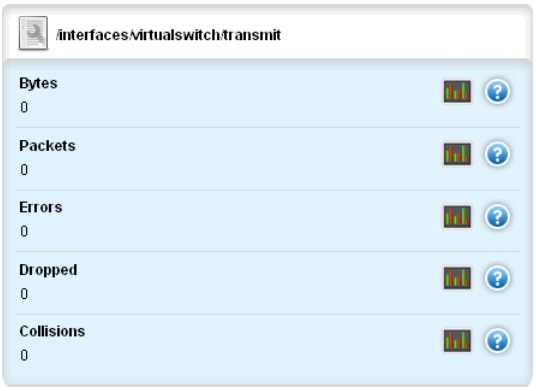

Figure 17.13. Transmit form

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets transmitted.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the transmit device.

Collisions

Synopsis: *unsigned integer*

Number of collisions detected on the port.

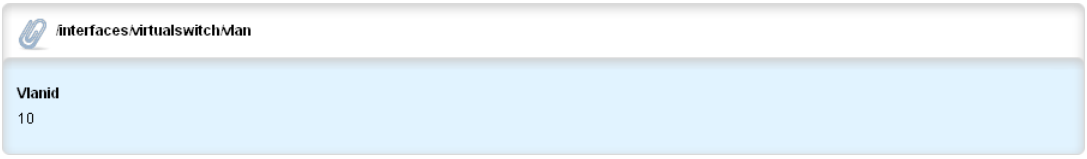

Figure 17.14. VLAN table

To display this table, navigate to *interfaces/virtualswitch/{virtualswitch number}/vlan*.

VLAN ID

Synopsis: *integer*

VLAN ID.

Figure 17.15. VLAN Receive form

To display the Receive and Transmit forms, navigate to *interfaces/virtualswitch/{virtualswitch number}/vlan/{number}*.

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped Into Abyss

Synopsis: *unsigned integer*

Number of dropped packets by the receive device.

Figure 17.16. VLAN Transmit form

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets transmitted.

Dropped

Synopsis: *unsigned integer*

Number of dropped packets by the transmit device.

Collisions

Synopsis: *unsigned integer*

Number of collisions detected on the port.

18. Modem

18.1. Modem Configuration and Status

The Modem menu is accessible from the main menu under *interface*. The path to this menu is *interface/modem*. The Routable Modem Interfaces table appears on the same screen as the Modem menu.

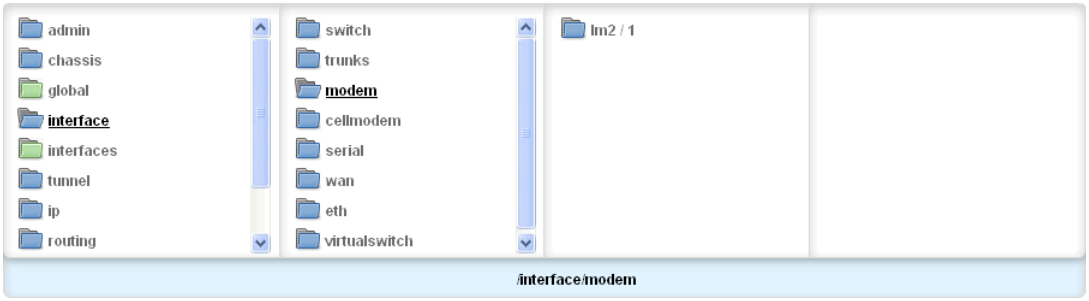

Figure 18.1. Modem menu

Figure 18.2. Routable Modem Interfaces table

The path to the Routable Modem Interfaces form is *interface/modem/{line module}*.

Figure 18.3. Routable Modem Interfaces form

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

enabled

Synopsis: *boolean*

Default: *true*

Provides the option to enable or disable this interface. When unchecked(i.e disabled), the interface will prevent all frames from being sent and received on that interface.

link-alarms

Synopsis: *boolean*

Default: *true*

Disabling link-alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that interface. Link alarms may also be controlled for the whole system under admin / alarm-cfg.

alias

Synopsis: *A string*

The SNMP alias name of the interface

Status information for the modem can be viewed by accessing the top-level *interfaces* menu. The path to the Interfaces Modem menu is *interfaces/modem*. The Modem Interfaces table appears on the same screen as the Interfaces Modem menu.

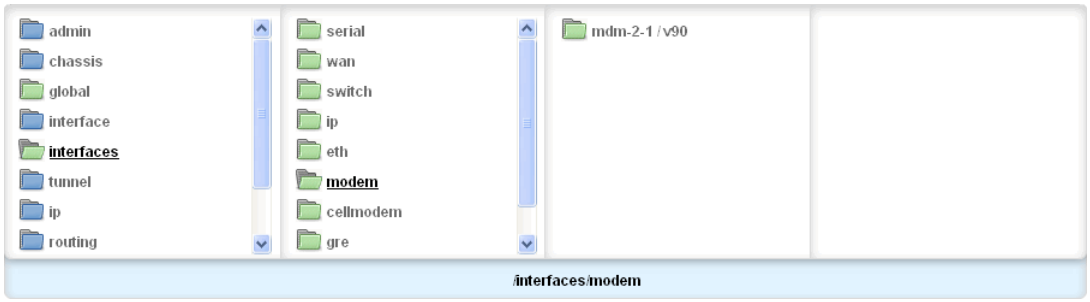

Figure 18.4. Interfaces Modem menu

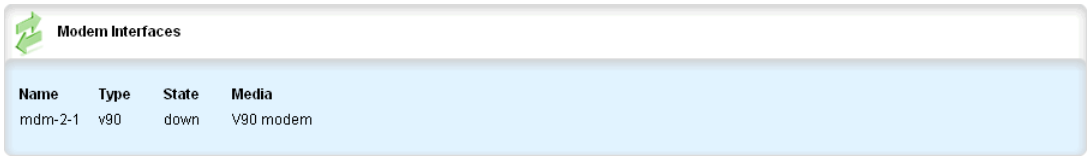

Figure 18.5. Modem Interfaces table

Name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

Interface name

type

Synopsis: *string - one of the following keywords { external, v90 }*

state

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

The port's link status.

media

Synopsis: *A string*

The type of port media { ***range of values** }. It provides the user with a description of the installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), and may be Short Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, the media description is displayed per the SFF-8472 specification, if the transceiver is plugged into the module. E.g. 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST,1000SX SFP LC S SL M5.

The path to the Interfaces Modem Submenu is *interfaces/modem* and then clicking on one of the linked submenus that follow (for example, mdm-2-1). The Modem Interfaces form, the PPP Interfaces Statistics form and the Reset and At trigger actions are accessible from the Interfaces Modem submenu.

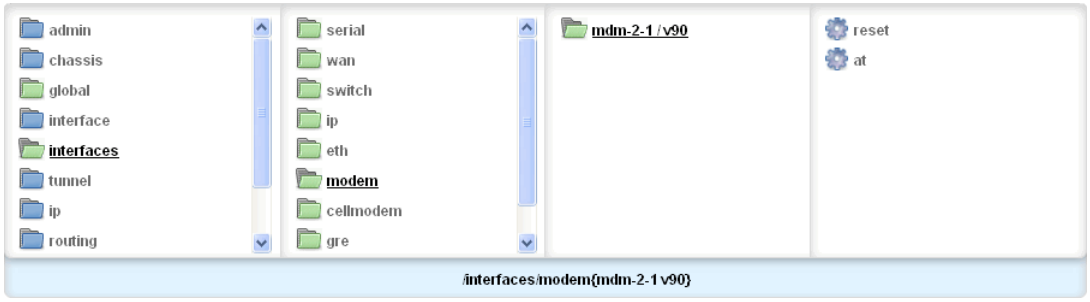

Figure 18.6. Interfaces Modem submenu

Figure 18.7. Modem Interfaces form

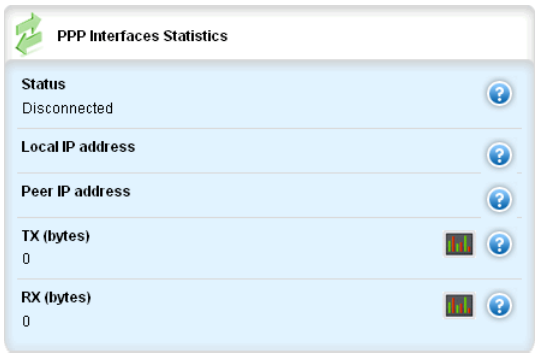

Figure 18.8. PPP Interfaces Statistics form

Status

Synopsis: *A string*

PPP connection status

Local IP address

Synopsis: *A string*

The IP address assigned to the modem by the remote server

Peer IP address

Synopsis: *A string*

The IP address of the remote server

TX (bytes)

Synopsis: *unsigned integer*

The bytes transmitted over the modem

RX (bytes)

Synopsis: *unsigned integer*

The bytes received by the modem

To reset the Modem, click on the *reset* action and then click the *Perform* button on the trigger action.

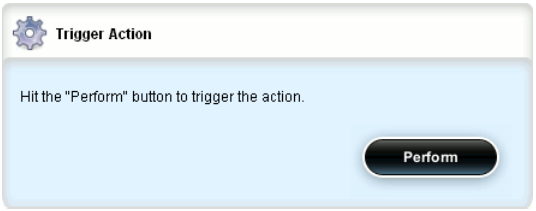

Figure 18.9. Reset Trigger Action

To launch the AT command, click on the *at* action, then enter a modem AT command into the *command* field of the AT command form and then click the *Perform* button on the trigger action.

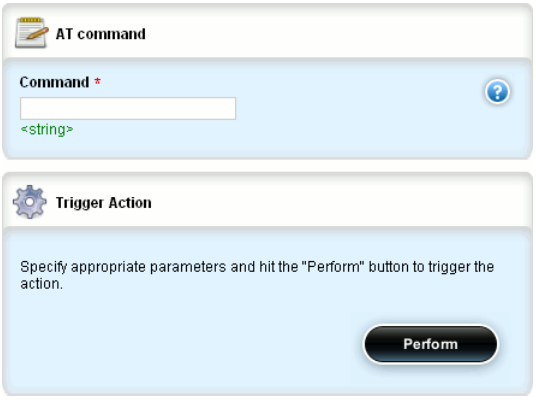

Figure 18.10. At Command and Trigger Action forms

If modems are configured, the Interface Modem submenu will be accessible from a link to the Modem menu, and will lead to forms used for configuring modem information for HSPA, Edge, CDMA, V90 and External. The path to this menu is *interface/modem* and then clicking on lm6 / 1 or any other linked submenus that are present. More information about the data accessible from this menu will be supplied later in this chapter.

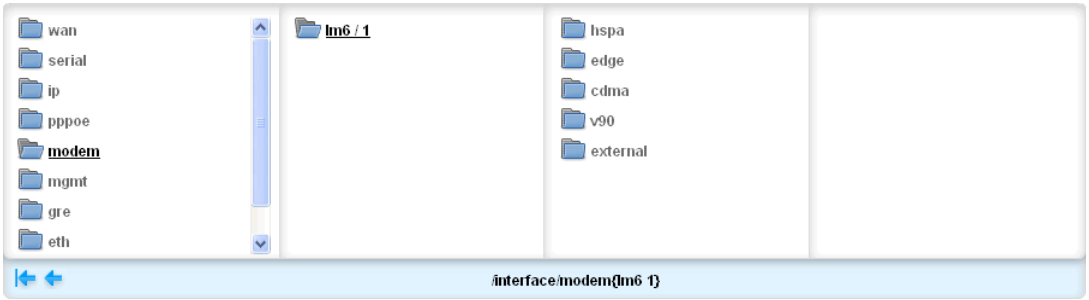

Figure 18.11. Interface Modem submenu

18.2. PPP and the Embedded Modem

18.2.1. PPP and Modem Fundamentals

RuggedRouter® may be equipped with an internal modem which allows connection to an external modem. A modem allows connections to be made over standard telephone lines. PPP (the Point-to-Point Protocol) is used to establish a network connection over a modem link.

18.2.1.1. PPP Interface

When a PPP connection is established, a network interface is created in the system. The interface name for both internal and external modem connections is *mdm-{slot}-1*, for example: mdm-2-1. Refer to this interface name when configuring firewall rules. A list of modem interfaces can be viewed at *ip/mdm-{slot}-1*. See Figure 18.12, “ip/mdm-slot-1” for an illustration.

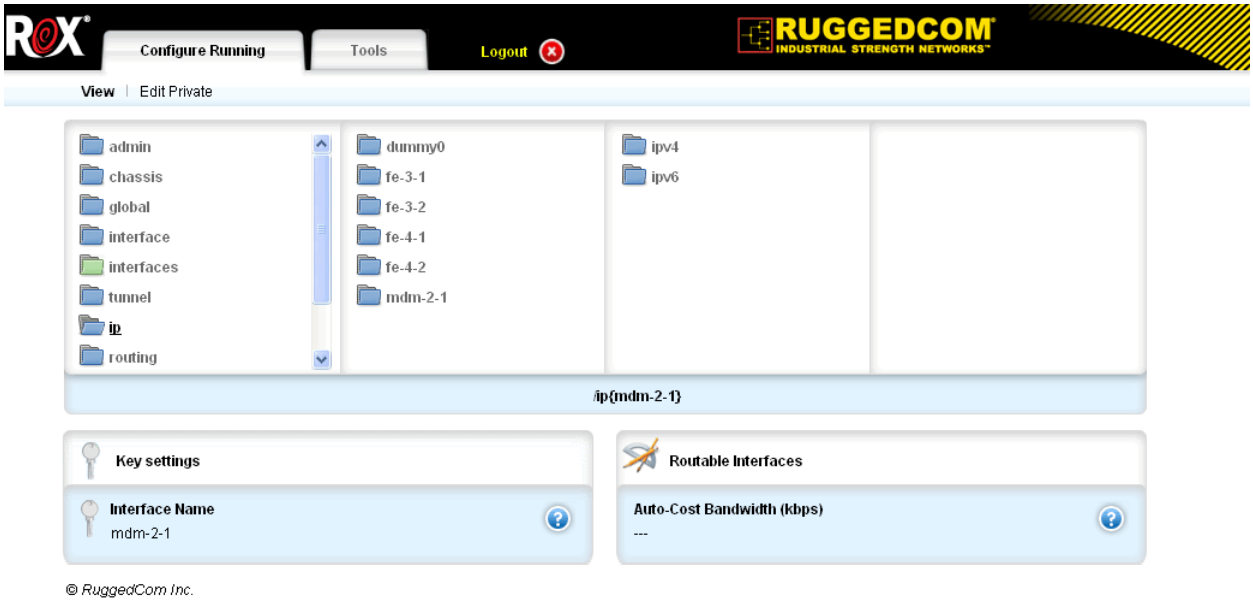

Figure 18.12. ip/mdm-slot-1

18.2.1.2. Authentication, Addresses and DNS Servers

PPP authentication will automatically use either the PAP or CHAP protocol.

In order to create a PPP client connection, you will need to obtain a user ID and password along with a telephone number from the operator of the PPP server that you will be dialing. The operator might be an Internet Service Provider or a system administrator within your organization.

The authentication process will provide a local IP address for use on the PPP interface and optionally the addresses of the DNS servers and a default gateway address to use. You should generally use these addresses unless you need to provide your own.

The PPP interface’s IP address, obtained from the PPP server, can be either a dynamic or a static IP address. Firewall configuration should be performed as is appropriate.

In the case of a PPP server configuration, you must configure the parameters described above for incoming PPP client connections.

18.2.1.3. When the Modem Connects

A PPP Client Connection may be configured to connect at boot time.

18.2.1.4. PPP Dial on Demand

The PPP client can be configured to dial only when there is traffic to be transmitted. In order to do that, the PPP interface must be configured to be the default gateway. On RuggedRouter®, the default gateway option is enabled automatically when the PPP client is configured to dial on demand. Initially, the PPP interface is shown as being up but the link is not actually connected. Only when there is traffic to be transmitted via the PPP interface is the configured number dialed and a connection initiated.

If dialing is successful and a PPP connection is established, the PPP server at the remote end of the connection will configure any remote IP addressing necessary.

18.2.1.5. V90

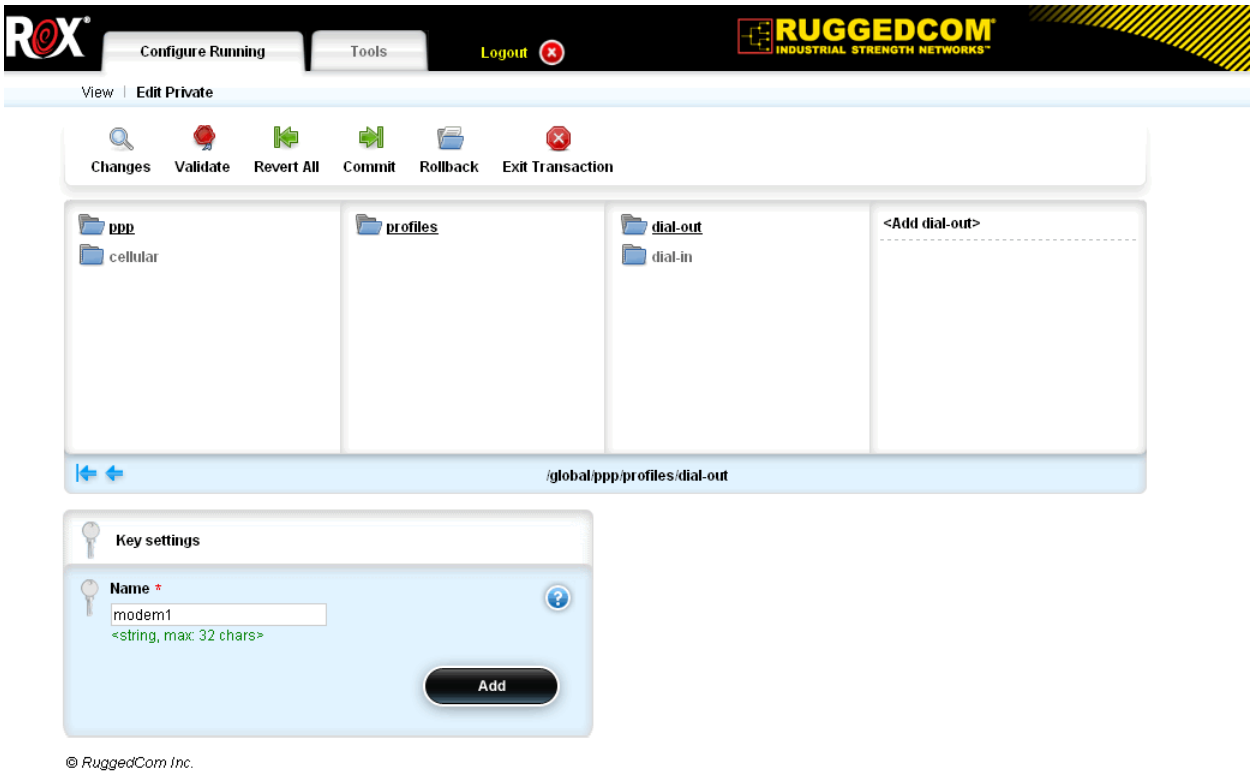

Figure 18.13. Configuring a PPP Profile for Dial-out

To configure a PPP profile for dial-out, enter Edit Private mode and go to the *global/ppp/profiles/dial-out* screen. Click on "Add dial-out". The Key Settings form will appear.

In the *Name* field of the Key Settings form, enter a name for your modem - for example, modem1. Then click the *Add* button. The PPP Configuration form will appear.

Figure 18.14. PPP Configuration form

Enter the required parameters into the PPP Configuration form.

Click the *Commit* button.

Click the *Exit Transaction* button.

The V90 Dial-out modem can be configured at *interface/modem/lm2/1/v90/ppp-client*. For more information, see Section 18.2.1.5.3, “Configuring V90 Modems”.

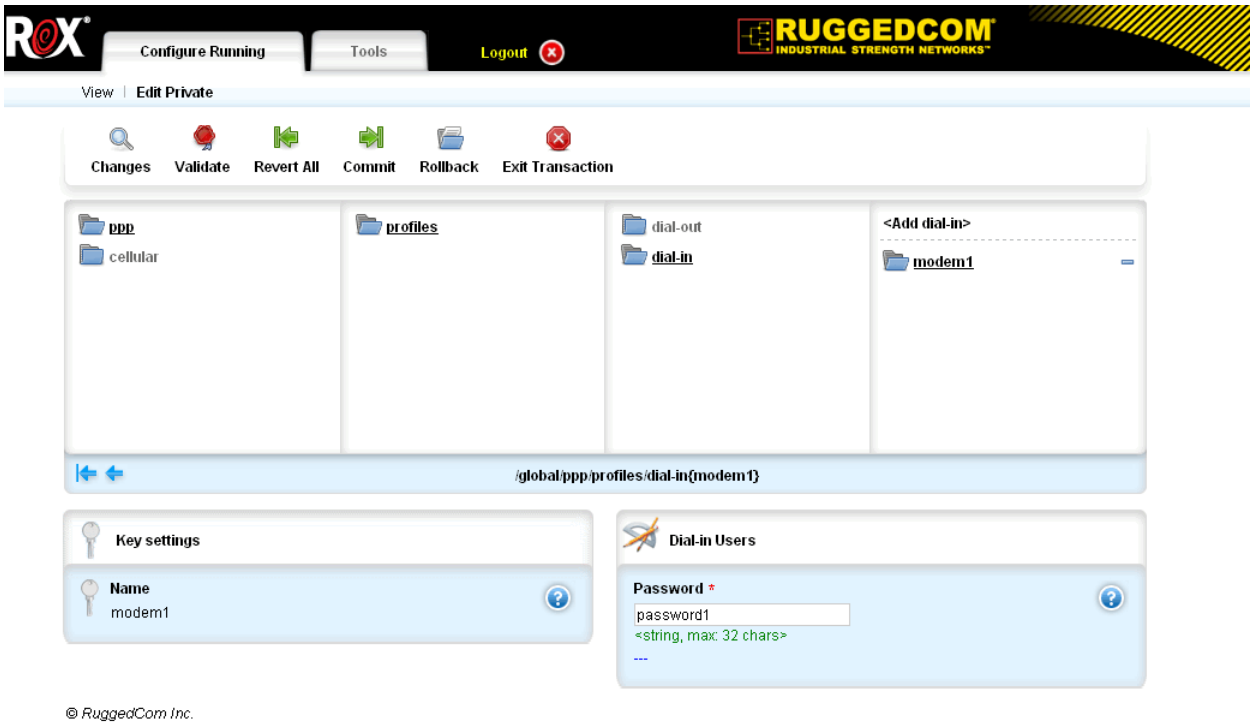

Figure 18.15. Configuring a PPP Profile for Dial-in

To configure a PPP profile for dial-in, enter Edit Private mode and go to the *global/ppp/profiles/dial-in* screen. Click on "Add dial-in". The Key Settings form will appear.

In the *Name* field of the Key Settings form, enter the modem name - for example, modem1. Then click the *Add* button. The Dial-in Users Form will appear.

In the *Password* field of the Dial-in Users form, enter the password you want to use - for example, password1.

Click the *Commit* button.

Click the *Exit Transaction* button.

The modem name and password will be saved. They can be viewed or modified in the Dial-in Users table by going to the *global/ppp/profiles/dial-in* screen. See Figure 18.16, “Dial-in PPP Users table” below.

Figure 18.16. Dial-in PPP Users table

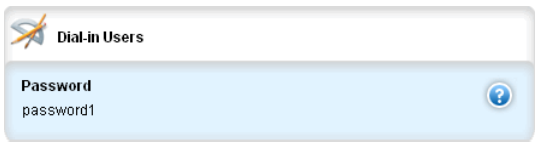

Figure 18.17. Dial-in Users form

As well, the password can be viewed or modified in the Dial-in Users form at *global/ppp/profiles/dial-in/modem{name}*.

The V90 Dial-in modem can be figured at *interface/modem/lm2/1/v90/ppp-server*. For more information on this procedure, see Section 18.2.1.5.3, “Configuring V90 Modems”.

18.2.1.5.3. Configuring V90 Modems

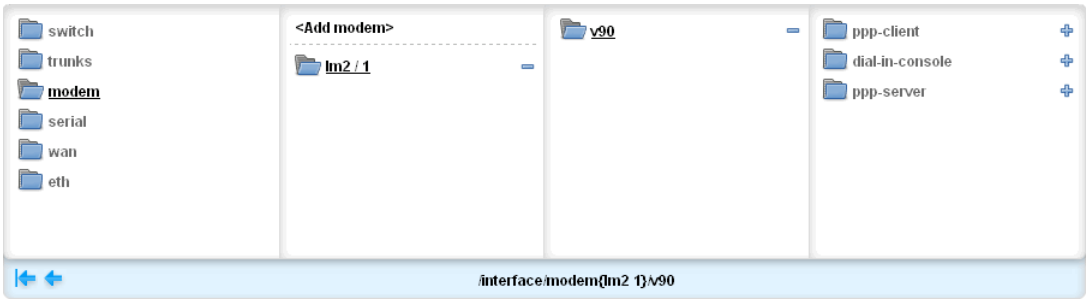

Figure 18.18. V90 menu

After PPP profiles have been created (see Section 18.2.1.5, “V90”), Dial-out and Dial-in modems can be configured from the V90 menu.

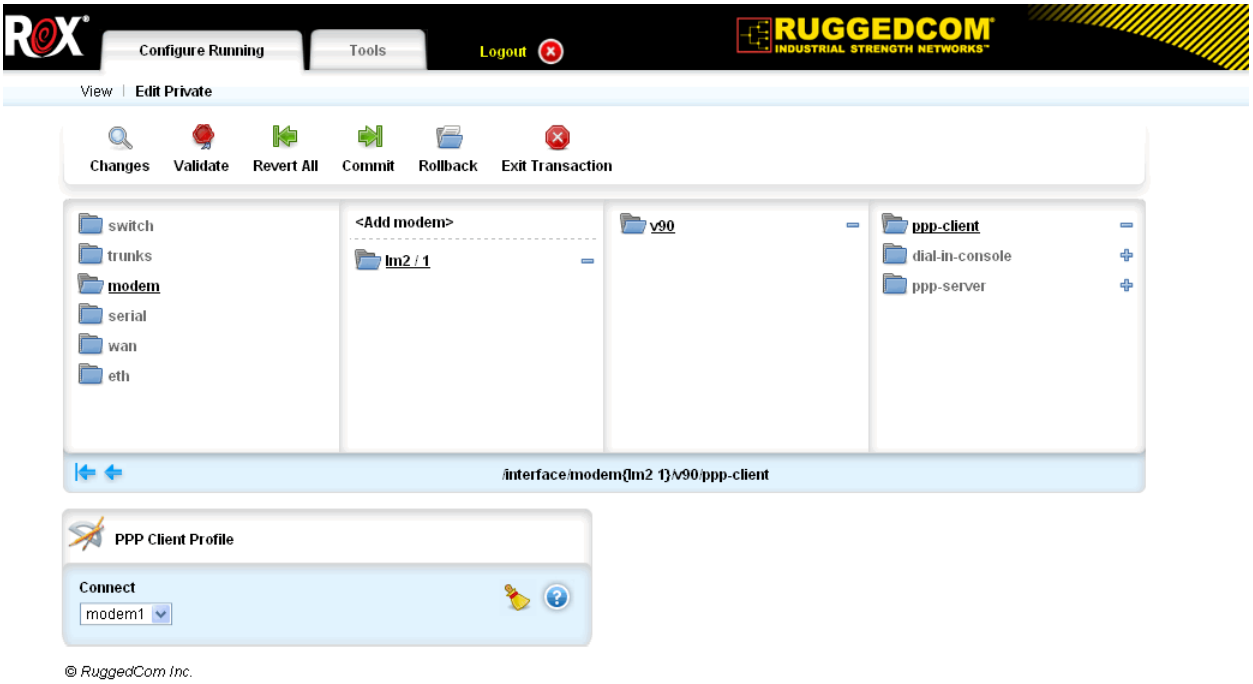

Figure 18.19. V90 PPP Client Configuration

To configure the Dial-out modem, click on the v90 submenu and then click on the plus sign icon next to the *ppp-client* submenu. The PPP Client Profile form will appear.

In the *Connect* field of the PPP Client Profile form, select the desired profile name from the list - for example, modem1.

Click the *Commit* button.

Click the *Exit Transaction* button or continue to configure the Dial-in modem, as shown below.

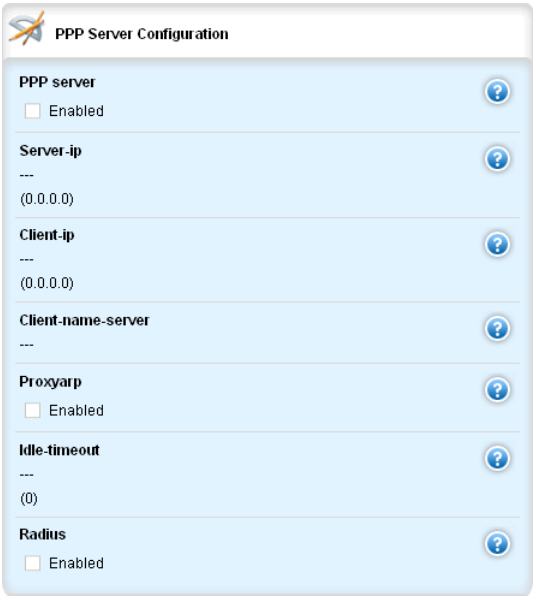

Figure 18.20. V90 PPP Server Configuration form

To configure the Dial-in modem, click on v90 submenu and then click on the plus sign icon next to the *ppp-server* submenu. The PPP Server Configuration form will appear.

Click "Enabled" in the PPP server field and fill in the required parameters (which are marked with an asterisk).

Click the *Commit* button.

Click the *Exit Transaction* button.

18.2.1.5.4. V90 Forms

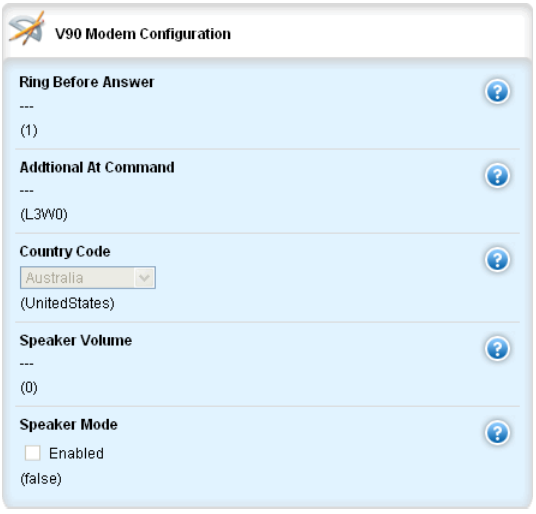

Figure 18.21. V90 Modem Configuration form

After modems have been configured, the path to the v90 Modem Configuration, Modem and PPP Server Configuration forms will be *interface/modem*, then clicking on a linked submenu (for example, lm6 / 1) and then clicking on *v90*.

ring-before-answer

Synopsis: *integer*

Default: *1*

The number of rings to wait before answering

addtional-at-command

Synopsis: *string*

Any extra AT command to use when initializing the modem

country-code

Synopsis: *string - one of the following keywords { Australia, Austria, Belgium, Brazil, China, Denmark, Finland, France, Germany, Greece, India, Ireland, Italy, Japan, Korea, Malaysia, Mexico, Netherlands, Norway, Poland, Portugual, SouthAfrica, Singapore, Spain, Sweden, Switzerland, TBR21, Taiwan, UnitedKingdom, UnitedStates, Venezuela }*

Default: *UnitedStates*

The modem country code

speaker-volume

Synopsis: *integer*

Default:

The modem speaker volume

speaker-mode

The modem speaker mode

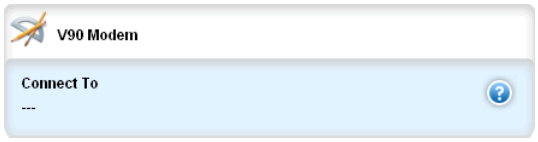

Figure 18.22. V90 Modem form

Connect

Synopsis: *A string*

Selects the PPP profile to start the PPP connection. The PPP profile is configured in /global/ppp/profiles/dial-out

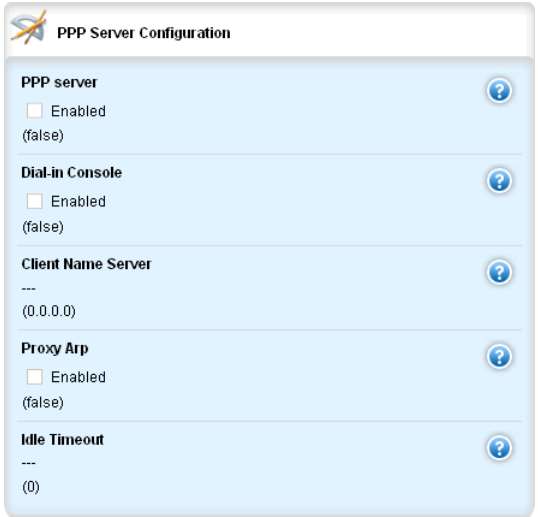

Figure 18.23. V90 PPP Server Configuration form

Configure the embedded PPP server for remote access via the modem.

PPP server

Enables the embedded PPP server to offer remote access via the modem

server-ip

Synopsis: *IPv4 address in dotted-decimal notation*

Default: *0.0.0.0*

The Local IP address of the PPP server

client-ip

Synopsis: *IPv4 address in dotted-decimal notation*

Default: *0.0.0.0*

The IP address to be assigned to the remote PPP endpoint

client-name-server

Synopsis: *IPv4 address in dotted-decimal notation*

The nameserver for the remote PPP client to look up

proxyarp

The option making the router attempt to proxy ARP the remote IP onto a local Ethernet subnet

idle-timeout

Synopsis: *integer*

Default:

The delay time before hanging up the PPP connection when there is no traffic.

Setting to 0 means disabling the idle timeout. The PPP connection will never be hung up.

radius

Enable radius server authentication for incomming PPP connection.

18.2.1.6. External

If data is configured, the path to the External Modem Configuration, Modem and PPP Server Configuration forms is *interface/modem/lm6 / 1/external*.

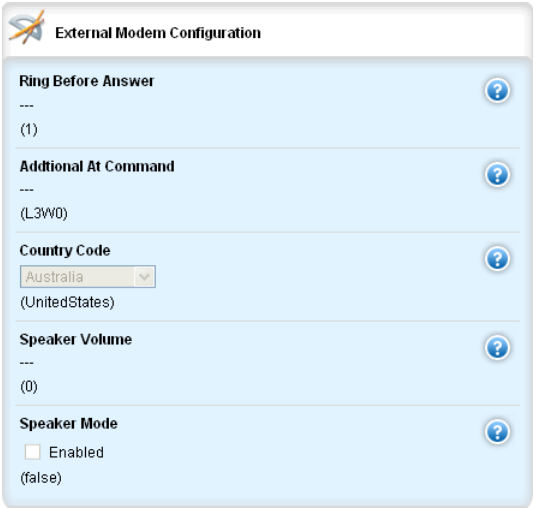

Figure 18.24. External Modem Configuration form

Configure the embedded modem.

ring-before-answer

Synopsis: *integer*

Default: *1*

The number of rings to wait before answering

addtional-at-command

Synopsis: *string*

Any extra AT command to use when initializing the modem

country-code

Synopsis: *string - one of the following keywords { Australia, Austria, Belgium, Brazil, China, Denmark, Finland, France, Germany, Greece, India, Ireland, Italy, Japan, Korea, Malaysia, Mexico, Netherlands, Norway, Poland, Portugual, SouthAfrica, Singapore, Spain, Sweden, Switzerland, TBR21, Taiwan, UnitedKingdom, UnitedStates, Venezuela }*

Default: *UnitedStates*

The modem country code

speaker-volume

Synopsis: *integer*

Default:

The modem speaker volume

speaker-mode

The modem speaker mode

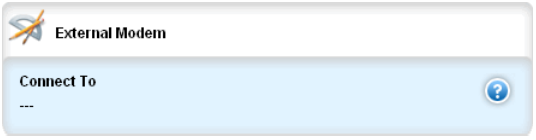

Figure 18.25. External Modem form

Connect

Synopsis: *A string*

Selects the PPP profile to start the PPP connection. The PPP profile is configured in /global/ppp/profiles/dial-out

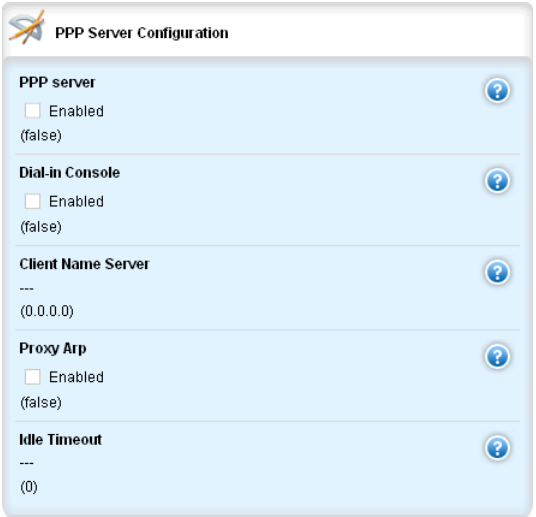

Figure 18.26. External PPP Server Configuration form

PPP server

Enables the embedded PPP server to offer remote access via the modem

server-ip

Synopsis: *IPv4 address in dotted-decimal notation*

Default: *0.0.0.0*

The Local IP address of the PPP server

client-ip

Synopsis: *IPv4 address in dotted-decimal notation*

Default: *0.0.0.0*

The IP address to be assigned to the remote PPP endpoint

client-name-server

Synopsis: *IPv4 address in dotted-decimal notation*

The nameserver for the remote PPP client to look up

proxyarp

The option making the router attempt to proxy ARP the remote IP onto a local Ethernet subnet

idle-timeout

Synopsis: *integer*

Default:

The delay time before hanging up the PPP connection when there is no traffic.

Setting to 0 means disabling the idle timeout. The PPP connection will never be hung up.

radius

Enable radius server authentication for incomming PPP connection.

18.3. PPP and the Cellular Modem

18.3.1. PPP and Cellular Modem Fundamentals

RX1000 may be equipped with an internal cellular modem or land-line modem. PPP (the Point-to-Point Protocol) is used to establish an IP network connection over a cellular radio modem link.

Depending on local cellular network availability, one of three cellular modem types may be ordered:

- Edge
- CDMA
- HSPA

18.3.1.1. PPP Interface

When a PPP connection is established, a network interface is created in the system. You can find the interface name . Refer to the interface name when configuring firewall rules.

18.3.1.2. Authentication, IP Addressing and DNS Servers

In contrast to the configuration for land-line modems, username and password might not be required for some cellular data service providers. If username and password is not required, you can enter none in the username and password fields of the GUI, or leave them blank. If authentication is required by the cellular data service provider, again PPP authentication will automatically use PAP or CHAP. Your service provider will provide you with a username and password along with an Access Provider Name (APN), which must be entered in the GUI.

The authentication process will provide a local IP address for use on the PPP interface and optionally the addresses of the DNS servers and a default gateway address to use. You should generally use these addresses unless you need to provide your own.

The PPP interface's IP address, obtained from the PPP server, can be either a dynamic or a static IP address. Firewall configuration should be performed as is appropriate.

18.3.1.3. When the Modem Connects

A PPP Client Connection for the cellular modem may be configured to connect at boot time.

18.3.2. PPP Cellular Modem Information and Configuration

The following sections review the forms used to view and configure HSPA, Edge, and CDMA cellular modems. The HSPA, Edge, and CDMA menus can be accessed from the *interfaces/cellmodem* menu below.

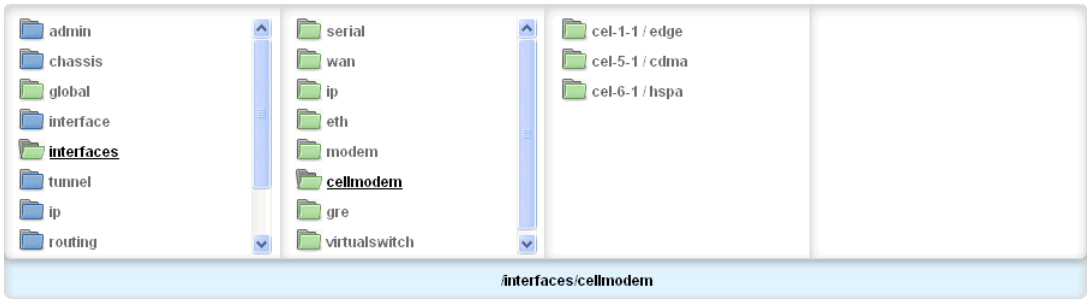

Figure 18.27. Interfaces Cellmodem menu

18.3.2.1. HSPA

The HSPA GSM profile is selected from the HSPA menu but Edge data needs to be configured from the Global Cellular GSM menu. See Section 18.3.2.3, "Global Cellular Modem GSM Configuration" for information on configuration.

If data is configured, the HSPA Cellular Modem Information form can be found under *interfaces/cellmodem/{line module}/hspa*.

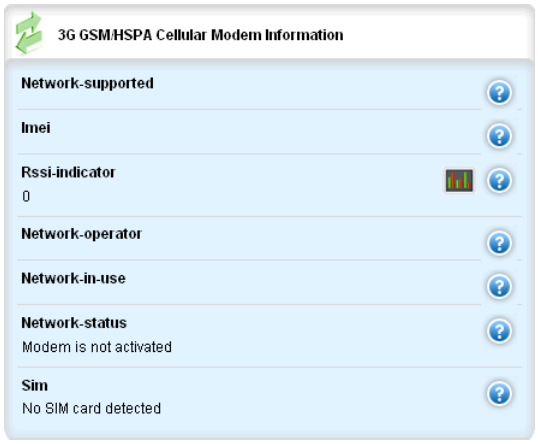

Figure 18.28. HSPA Cellular Modem Information form

The HSPA Cellular Modem Information form displays modem information.

network-supported

Synopsis: *A string*

Wireless technologies supported by the modem

imei

Synopsis: *A string*

International Mobile Equipment Indentity

radio

Synopsis: *A string*

The current RF status of cellmodem

rssi-indicator

Synopsis: *integer*

The Received Signal Strength Indicator in dBm

network-operator

Synopsis: *A string*

The wireless network operator currently in use

network-in-use

Synopsis: *A string*

The network technology currently in use by the modem

network-status

Synopsis: *A string*

The registration status of the modem with the wireless network

sim

Synopsis: *A string*

The Subscriber Indentity Module number

The following information provides additional details about the fields in the HSPA Cellular Modem Information Form.

The *IMEI* (International Mobile Equipment Identity) is a numeric identifier unique to the cellular modem card.

Rssi Indicator (Received Signal Strength) indicates the signal level received by the cellular modem from the cell site.

Network Operator displays the identity of the wireless network provider to which the cellular modem is currently connected.

The *Network In Use* field displays which network technology is currently in use between the modem and the network.

Network Status displays the current registration status of the cellular modem with respect to the cellular network. Possible values are:

- Registered, home
- Registered, roaming
- Unregistered

SIM displays the ID of the SIM card currently installed in the cellular modem.

18.3.2.2. Edge

The Edge GSM profile is selected from the Edge menu but Edge data needs to be configured from the Global Cellular GSM menu. See Section 18.3.2.3, “Global Cellular Modem GSM Configuration” for information on configuration.

If data is configured, the path to the Edge Cellular Modem Information form is *interfaces/cellmodem/{line module}/edge*.

Figure 18.29. Edge Cellular Modem Information form

network-supported

Synopsis: *A string*

Wireless technologies supported by the modem

imei

Synopsis: *A string*

International Mobile Equipment Indentity

radio

Synopsis: *A string*

The current RF status of cellmodem

rssi-indicator

Synopsis: *integer*

The Received Signal Strength Indicator in dBm

network-operator

Synopsis: *A string*

The wireless network operator currently in use

network-in-use

Synopsis: *A string*

The network technology currently in use by the modem

network-status

Synopsis: *A string*

The registration status of the modem with the wireless network

sim

Synopsis: *A string*

The Subscriber Indentity Module number

The following information provides additional details about the fields in the Edge Cellular Modem Information Form.

Rssi Indicator (Received Signal Strength) indicates the signal level received by the cellular modem from the cell site.

Network Operator displays the identity of the wireless network provider to which the cellular modem is currently connected.

Network Status displays the current registration status of the cellular modem with respect to the cellular network. Possible values are:

- Registered, home
- Registered, roaming
- Unregistered

SIM displays the ID of the SIM card currently installed in the cellular modem.

18.3.2.3. Global Cellular Modem GSM Configuration

Figure 18.30. Global Cellular GSM menu

HSPA and Edge data is configured from the Global Cellular GSM menu.

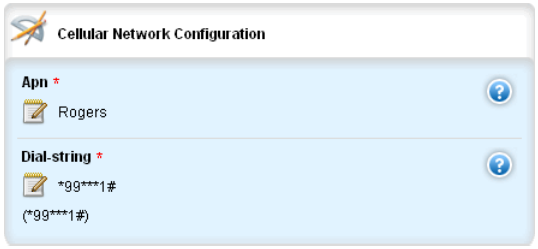

Figure 18.31. GSM Cellular Network Configuration form

name

Synopsis: *A string*

Create gsm profile name

apn

Synopsis: *string*

The wireless network access point name

dial-string

Synopsis: *string*

Default: **99***1#*

The dial string given by the wireless provider to connect to the access point name

The Access Point Name (APN) is necessary only on GPRS networks (Edge or HSPA). It is the name of the cellular network access point which provides a gateway to the Internet. This information will be

provided by the wireless network when you register for data service. This field is not used for CDMA modems.

The Dial string is a special command to be sent by the cellular modem to the cellular network to establish a data connection. For example, for GSM/GPRS networks, this is typically: *99***1#. This command will depend on the wireless network. Please consult the wireless network operator for the correct dial string command for data service. A regular telephone number is usually not required to connect to a GSM/GPRS network.

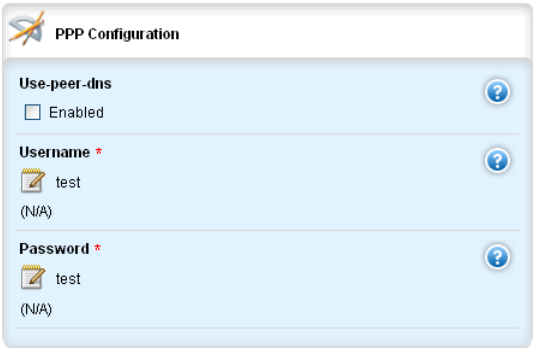

Figure 18.32. PPP Configuration form

use-peer-dns

Enables the DNS server entries that the PPP server recommends. Enables this option unless you provide your own name servers

username

Synopsis: *string*

Default: *N/A*

The user ID to connect to the remote server

password

Synopsis: *string*

Default: *N/A*

The password to be authenticated by the remote server

dial-on-demand

Activates Dial-on-Demand on this connection. The establishment of the PPP connection is postponed until there is data to be transmitted via the interface

disconnect-idle-timeout

Synopsis: *integer*

Default:

The time in seconds to wait before disconnecting PPP when there is no traffic on the link. This option is only valid when Dial-on-Demand is enabled

failover-on-demand

Activate link failover on-demand on this device. PPP link establishment on this device is controlled by link failover

18.3.2.4. CDMA

The CDMA GSM profile is selected by using the CDMA EVDO Cellular Modem Configuration form but the profile needs to be configured first from the Global Cellular CDMA Modem menu. See Section 18.3.2.4.2, “Global Cellular CDMA Modem Configuration” for information on configuration. After

the profile has been configured, the CDMA functions and the CDMA EVDO modem Configuration form are accessible from the CDMA menu.

The path to the CDMA menu is *interfaces/cellmodem/{line module}/cdma*. The CDMA EVDO Cellular Modem Information form appears on the same screen as the *CDMA* menu.

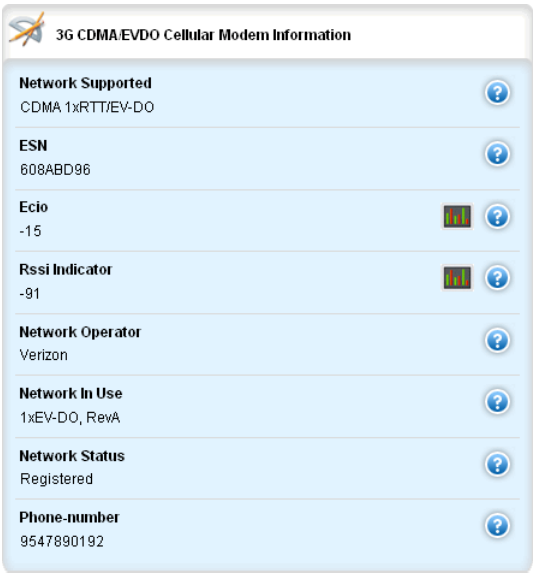

Figure 18.33. CDMA EVDO Cellular Modem Information form

network-supported

Synopsis: *A string*

Wireless technologies supported by the modem

esn

Synopsis: *A string*

The Electronic Serial Number of the modem. ESN is only avaible for the CDMA modem.

ecio

Synopsis: *integer*

The total energy per chip per power density value in dBm

rssi-indicator

Synopsis: *integer*

The Received Signal Strength Indicator in dBm

network-operator

Synopsis: *A string*

The wireless network operator currently in use

network-in-use

Synopsis: *A string*

The network technology currently in use by the modem

network-status

Synopsis: *A string*

The registration status of the modem with the wireless network

phone-number

Synopsis: *A string*

The subscriber phone number of the CDMA modem

The information below provides additional details about the fields in the CDMA EVDO Cellular Modem Information form.

The *Electronic Serial Number (ESN)* is a numeric identifier unique to the cellular modem card. This corresponds to the IMEI for GSM networks.

Rssi Indicator (Received Signal Strength) indicates the signal level received by the cellular modem from the cell site.

Network Operator displays the identity of the wireless network provider to which the cellular modem is currently connected.

The *Network In Use* field displays which network technology is currently in use between the modem and the network.

Network Status displays the current registration status of the cellular modem with respect to the cellular network. Possible values are:

- Registered, home
- Registered, roaming
- Unregistered

Phone Number displays the cellular telephone number associated with the account created to provide service for the modem.

18.3.2.4.1. Cellular Modem Account Activation

Prior to use, a CDMA-type cellular modem must be activated for use on a particular provider's network. Once the activation process has been completed, the modem will be able to connect to the network without further intervention. Two account activation methods are provided by ROX™: "OTA (Over-the-Air)" and "Manual". Both activation methods are described in this section.

Over-The-Air Account Activation

ROX™ supports the OTASP (Over-the-Air Service Provisioning) mechanism offered by most CDMA cellular service providers for provisioning cellular end stations for use on their networks. Using this method, the service provider, or carrier, supplies an OTASP dial string which ROX™ can use to contact the cellular network via the modem. During this OTASP call, the carrier authorizes and configures the modem for use on its network. Note that an OTASP dial string typically begins with "*228".

The path to the CDMA Over The Air Activation form and Trigger Action form is *interface/modem/lm6 / 1/cdma/OverTheAirActivation*.

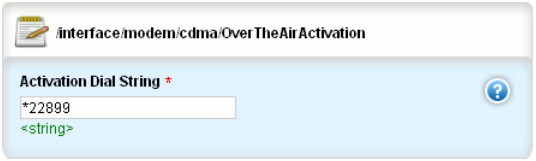

Figure 18.34. CDMA Over The Air Activation form

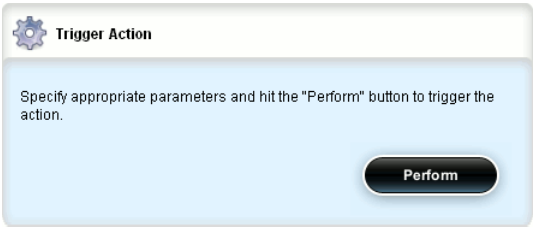

Figure 18.35. CDMA Over The Air Activation Trigger Action form

1. First, establish an account with the help of a service representative of the cellular network provider.
2. Enter the OTASP dial string supplied in the *Activation Dial string* field of the Over The Air Activation form.
3. Click the *Perform* button on the *Trigger Action* form.

Manual Account Activation

If the carrier does not support Over-the-Air Service Provisioning, the cellular modem must be programmed via the *Manual Account Activation* form using settings supplied by the carrier's service personnel:

The path to the CDMA Manual Activation form and Trigger Action form is *interface/modem/lm6 / 1/ cdma/ManualActivation*.

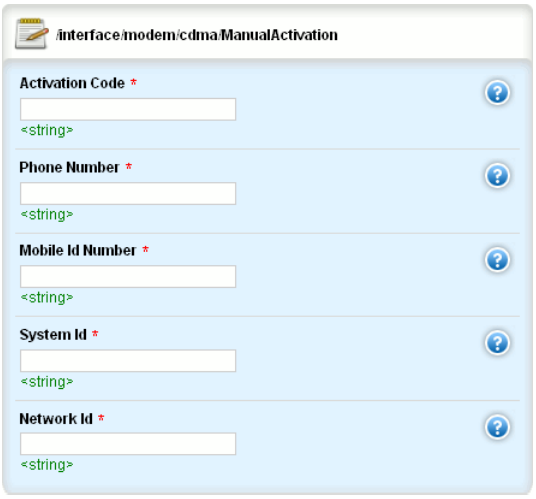

Figure 18.36. CDMA Manual Activation form

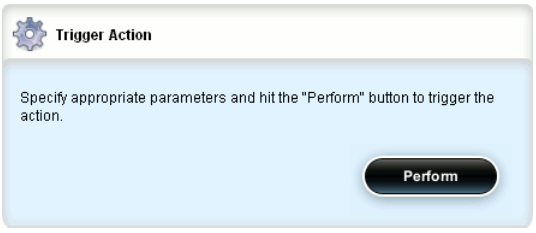

Figure 18.37. CDMA Manual Activation Trigger Action form

1. First, establish an account with a service representative of the cellular network provider. You will need the following settings in order to activate your modem. Note that not all of these parameters are required by all network providers:
   - *Activation code*, also known as a "subsidy lock".

- *Phone Number*, or MDN (Mobile Directory Number).
- *MIN* (Mobile Identification Number), often the same as the Phone Number.
- *System ID*, or Home System ID.
- *Network ID*

2. After specifying the parameters above in the Manual Activation form, click the *Perform* button on the *Trigger Action* form.

Reset Modem

The modem can be reset using the Reset Modem Trigger Action form. The path to the CDMA Reset Modem Trigger Action form is *interface/modem/lm6 / 1/cdma/resetmodem*.

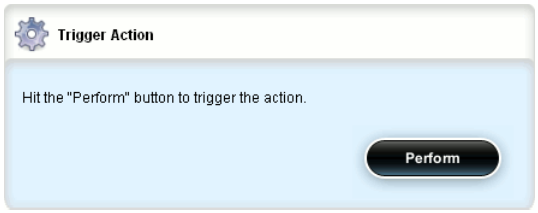

Figure 18.38. CDMA Reset Modem Trigger Action form

18.3.2.4.2. Global Cellular CDMA Modem Configuration

Figure 18.39. Global Cellular CDMA menu

The path to this menu is *global/cellular/profiles/cdma*. The Cellular Network Configuration table appears on the same screen as the *global* menu.

CDMA data is configured from the Global Cellular CDMA menu.

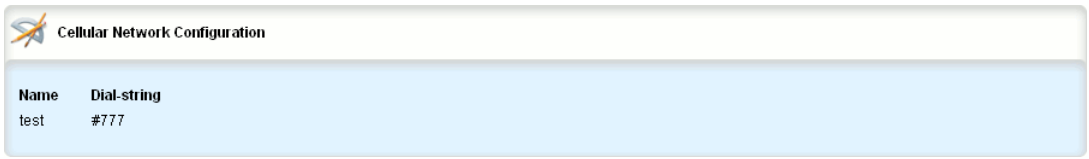

Figure 18.40. Cellular Network Configuration table

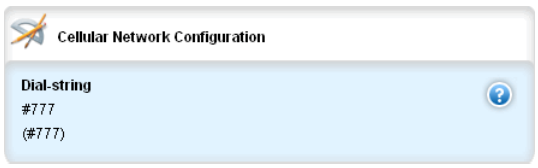

Figure 18.41. Cellular Network Configuration form

The Cellular Network Configuration form and the PPP Configuration form appear on the same screen as the *global* menu.

name

Synopsis: *A string*

Create cdma profile name

dial-string

Synopsis: *A string*

Default: *#777*

The dial string to connect to the wireless provider

The Dial string is a special command to be sent by the cellular modem to the cellular network to establish a data connection. For example, for GSM/GPRS networks, this is typically: *99***1#. This command will depend on the wireless network. Please consult the wireless network operator for the correct dial string command for data service. A regular telephone number is usually not required to connect to a GSM/GPRS network.

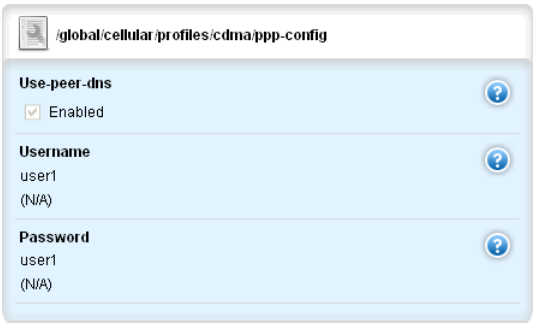

Figure 18.42. PPP Configuration form

use-peer-dns

Enables the DNS server entries that the PPP server recommends. Enables this option unless you provide your own name servers

username

Synopsis: *string*

Default: *N/A*

The user ID to connect to the remote server

password

Synopsis: *string*

Default: *N/A*

The password to be authenticated by the remote server

dial-on-demand

Activates Dial-on-Demand on this connection. The establishment of the PPP connection is postponed until there is data to be transmitted via the interface

disconnect-idle-timeout

Synopsis: *integer*

Default:

The time in seconds to wait before disconnecting PPP when there is no traffic on the link. This option is only valid when Dial-on-Demand is enabled

failover-on-demand

Activate link failover on-demand on this device. PPP link establishment on this device is controlled by link failover

18.3.2.5. CellModem

18.3.2.5.1. CellModem Configuration

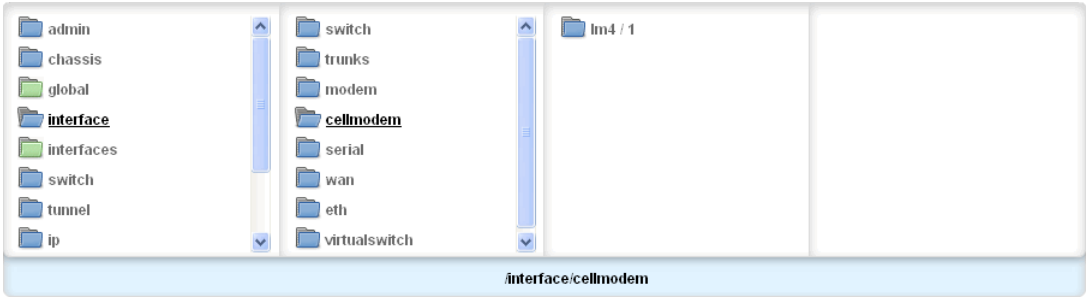

Figure 18.43. Interface Cellmodem menu

The path to the *interface/cellmodem* menu is *interface/cellmodem*. The Routable Cellular Modem Interfaces table appears on the same screen as this menu.

Figure 18.44. Routable Cellular Modem Interfaces table

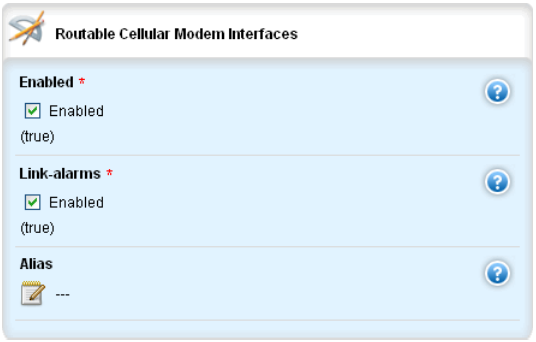

Figure 18.45. Routable Cellular Modem Interfaces form

The path to this form is *interface/cellmodem/{line module}*.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

enabled

Synopsis: *boolean*

Default: *true*

Provides the option to enable or disable this interface. When unchecked(i.e disabled), the interface will prevent all frames from being sent and received on that interface.

link-alarms

Synopsis: *boolean*

Default: *true*

Disabling link-alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that interface. Link alarms may also be controlled for the whole system under admin / alarm-cfg.

alias

Synopsis: *A string*

The SNMP alias name of the interface

Figure 18.46. Interface Cellmodem HSPA menu

The path to this menu is *interface/cellmodem/{line module}/hspa*.

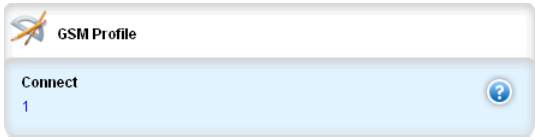

Figure 18.47. GSM Profile form

The path to this form is *interface/cellmodem/{line module}/hspa/ppp-client*.

Connect

Synopsis: *A string*

Selects the gsm profile to connect to wireless network. The gsm profile is configured in /global/cellular/profiles/gsm

18.3.2.5.2. CellModem Status

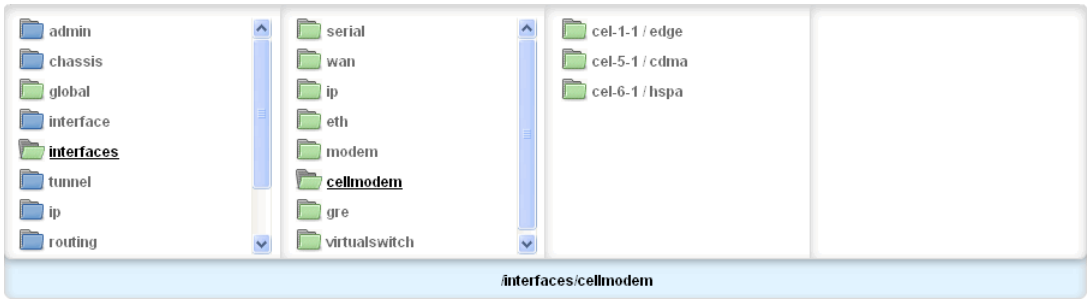

Figure 18.48. Interfaces Cellmodem menu

The path to the *interfaces/cellmodem* menu is *interfaces/cellmodem*. The Cellular Modem Interfaces table appears on the same screen as this menu.

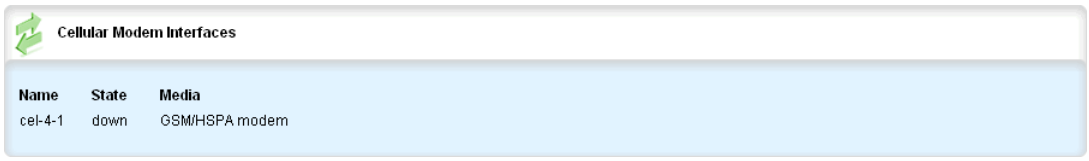

Figure 18.49. Cellular Modem Interfaces table

Name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

Interface name

state

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

The port's link status.

media

Synopsis: *A string*

The type of port media { ***range of values** }. It provides the user with a description of the installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), and may be Short Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, the media description is displayed per the SFF-8472 specification, if the transceiver is plugged into the module. E.g. 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST,1000SX SFP LC S SL M5.

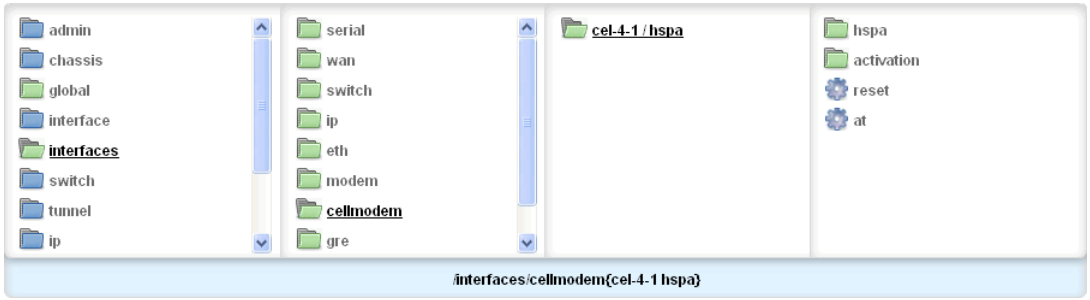

Figure 18.50. Interfaces Cellmodem HSPA menu

The path to this menu is *interfaces/cellmodem/{line module}/hspa*. The Cellular Modem Interfaces form and the HSPA PPP Interfaces Statistics form appear on the same screen as this menu.

Figure 18.51. Cellular Modem Interfaces form

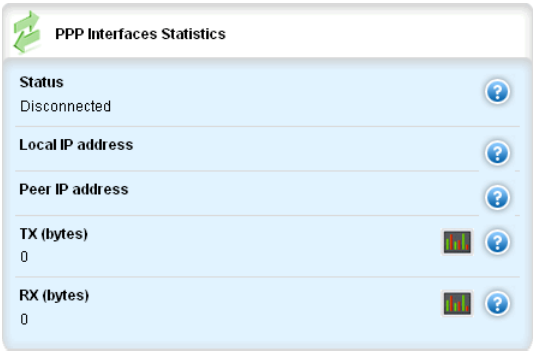

Figure 18.52. HSPA PPP Interfaces Statistics form

Status

Synopsis: *A string*

PPP connection status

Local IP address

Synopsis: *A string*

The IP address assigned to the modem by the remote server

Peer IP address

Synopsis: *A string*

The IP address of the remote server

TX (bytes)

Synopsis: *unsigned integer*

The bytes transmitted over the modem

RX (bytes)

Synopsis: *unsigned integer*

The bytes received by the modem

19. Serial Protocols

19.1. Introduction

This chapter familiarizes the user with:

- RawSocket Applications
- TCP Modbus Server Applications
- DNP (Distributed Network Protocol)
- Configuring Serial ports for RawSocket
- Viewing Serial Port and TCP Connection Status and Statistics
- Resetting Serial ports
- Tracing Serial Port activity

19.1.1. Serial IP Port Features

The RuggedCom Serial Server provides the following features for forwarding serial traffic over IP:

- Raw Socket Protocol - a means to transport streams of characters from one serial port on the router to a specified remote IP address and TCP port
- RX1000 supports 4 or 8 serial ports
- Bit rates of 300, 600, 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200, or 230400 bps.
- Supports RS232, RS422 and RS485 party line operation.
- XON/XOFF flow control.
- Supports a point-to-point connection mode and a broadcast connection mode in which up to 32 remote servers may connect to a central server.
- TCP/IP incoming, outgoing or both incoming/outgoing connections mode, configurable local and remote TCP port numbers.
- Packetize and send data on a full packet, a specific character or upon a timeout.
- Supports a “turnaround” time to enforce minimum times between messages sent out the serial port.
- Debugging facilities including connection tracing and statistics

19.1.1.1. LED Designations

Each RJ45 connector on the S01 serial module features two LEDs. When the serial card is located in the top slot, the transmit LED is to the left of the port, and the receive LED is to the right of the port. The transmit LED illuminates when characters are transmitted; the receive LED illuminates when characters are received.

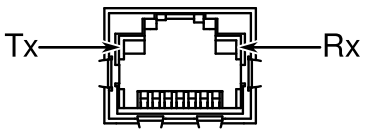

Figure 19.1. S01 Serial Module RJ45 Connector LEDs

19.1.2. Serial Protocols Applications

19.1.2.1. Character Encapsulation

Character encapsulation is used any time a stream of characters must be reliably transported across a network.

The character streams can be created by any serial device. The baud rates supported at either server need not be the same. If configured, the router will obey XON/XOFF flow control from the end devices.

One of the servers is configured to listen to TCP connection requests on a specific TCP port number. The other server is configured to connect to its peer on the listening port number. The RX1000 will periodically attempt to connect if the first attempt fails and after a connection is broken.

The RX1000 can be used to connect to any device supporting TCP (e.g. a host computer's TCP stack or a serial application on a host using port redirection software).

19.1.2.2. RTU Polling

The following applies to a variety of RTU protocols besides ModBus RTU, including ModBus ASCII and DNP.

The remote router communicates with host equipment through:

- native TCP connections,
- another RX1000 via a serial port or
- a port redirection package which Supports TCP.

If a RX1000 is used at the host end, it will wait for a request from the host, encapsulate it in a TCP message, and send it to the remote side. There, the remote RX1000 will forward the original request to the RTU. When the RTU replies, the the RX1000 will forward the encapsulated reply back to the host end.

ModBus does not employ flow-control so XON/XOFF should not be configured.

The RX1000 maintains configurable timers to help decide replies and requests are complete and to handle special messages such as broadcasts.

The RX1000 will also handle the process of line-turnaround when used with RS485.

19.1.2.3. Broadcast RTU Polling

Broadcast polling allows a single host connected RX1000 to "fan-out" a polling stream to a number of remote RTUs.

The host equipment connects via a serial port to a RX1000. Up to 32 remote devices may connect to the host server via the network.

Initially, the remote servers will place connections to the host server. The host server in turn is configured to accept the required number of incoming connections.

The host will sequentially poll each RTU. Each poll received by the host server is forwarded (i.e. broadcast) to all of the remote servers. All RTUs will receive the request and the appropriate RTU will issue a reply. The reply is returned to the host server, where it is forwarded to the host.

19.1.3. Serial Protocols Concepts And Issues

19.1.3.1. Host And Remote Roles

The RX1000 can either initiate or accept a TCP connection for serial encapsulation. It can establish a connection from field (“remote”) equipment to the central site (“host”) equipment, vice versa, or bi-directionally.

Configure the RX1000 at the host end to connect to the remote when:

- The host end uses a port redirector that must make the connection.
- The host end is only occasionally activated and will make the connection when it becomes active.
- A host end firewall requires the connection to be made outbound.

Connect from the remote to the host if the host end accepts multiple connections from remote ends in order to implement broadcast polling.

Connect from each side to other if both sides support this functionality.

19.1.3.2. Use Of Port Redirectors

Port redirectors are PC packages that emulate the existence of communications ports. The redirector software creates and makes available these “virtual” COM ports, providing access to the network via a TCP connection.

When a software package uses one of the virtual COM ports, a TCP connection is placed to a remote IP address and TCP port that has been programmed into the redirector. Some redirectors also offer the ability to receive connections.

19.1.3.3. Message Packetization

The server buffers received characters into packets in order to improve network efficiency and demarcate messages.

The server uses three methods to decide when to packetize and forward the buffered characters to the network:

- Packetize on Specific Character,
- Packetize on timeout and
- Packetize on full packet.

If configured to packetize on a specific character, the server will examine each received character and will packetize and forward upon receiving the specific character. The character is usually a <CR> or an <LF> character but may be any ASCII character.

If configured to packetize on a timeout, the server will wait for a configurable time after receiving a character before packetizing and forwarding. If another character arrives during the waiting interval, the timer is restarted. This method allows characters transmitted as part of an entire message to be forwarded to network in a single packet, when the timer expires after receiving the very last character of the message. This is usually the only packetizer selected when supporting ModBus communications.

Finally, the server will always packetize and forward on a full packet, i.e. when the number of characters fills its communications buffer (1024 bytes).

19.1.3.4. Use of Turnaround Delays

Some RTU protocols (such as ModBus) use the concept of a turnaround delay. When the host sends a message (such as a broadcast) that does not invoke an RTU response, it waits a turnaround delay

time. This delay ensures that the RTU has time to process the broadcast message before it has to receive the next poll.

When polling is performed, network delays may cause the broadcast and next poll to arrive at the remote server at the same time. Configuring a turnaround delay will enforce a minimum separation time between each message sent out the port. Note that turnaround delays do not need to be configured at the host computer side and may be disabled there.

19.1.4. TcpModBus Server Application

The TcpModbus Server application is used to transport Modbus requests and responses across IP networks.

The source of the polls is a Modbus “master”, a host computer that issues the polls over a serial line.

A TcpModbus Client application, such as that implemented by the RuggedServer accepts Modbus polls on a serial line from a master and determines the IP address of the corresponding RTU. The client then encapsulates the message in TCP and forwards the frame to a Server Gateway or native TcpModbus RTU. Returning responses are stripped of their TCP headers and issued to the master.

The TcpModbus Server application accepts TCP encapsulated modbus messages from Client Gateways and native masters. After removing the TCP headers the messages are issued to the RTU. Responses are TCP encapsulated and returned to the originator.

A “native” TcpModbus master is one that can encapsulate the Modbus polls in TCP and directly issue them to the network.

19.1.4.1. Local Routing At The Server Gateway

The Server Gateway supports up to 32 RTUs on any of its four ports. When a request for a specific RTU arrives the server will route it to the correct port.

19.1.4.2. MultiMaster Capability

It is possible for multiple masters to simultaneously issue requests for the same RTU. The Server Gateway will queue the requests and deliver them to the RTU in turn. This “multimaster” capability allows widely distributed masters to configure and extract information from the RTU.

19.1.5. TcpModbus Concepts And Issues

19.1.5.1. Host And Remote Roles

Client gateways (such as that implemented by the RX1000) always make the TCP connection to the Server Gateway. The Server Gateway can only accept a connection.

19.1.5.2. Port Numbers

The TCP port number dedicated to Modbus use is port 502. The Server Gateway can also be configured to accept a connection on a configurable port number. This auxiliary port can be used by masters that do not support port 502.

The Server Gateway is capable of creating only one connection on the specified auxiliary port, whereas when Modbus is configured to use the default port, 502, it may connect to multiple RTUs.

19.1.5.3. Retransmissions

The Server Gateway offers the ability to resend a request to an RTU should the RTU receive the request in error or the Server Gateway receives the RTU response in error.

The decision to use retransmissions, and the number to use depends upon factors such as:

- The probability of a line failure
- The number of RTUs and amount of traffic on the port
- The cost of retransmitting the request from the server vs. timing-out and retransmitting at the master. This cost is affected by the speed of the ports and of the network.

19.1.5.4. ModBus Exception Handling

If the Server Gateway receives a request for an unconfigured RTU, it will respond to the originator with a special message called an exception (type 10). A type 11 exception is returned by the server if the RTU fails to respond to requests.

Native TcpModbus polling packages will want to receive these messages. Immediate indication of a failure can accelerate recovery sequences and reduce the need for long timeouts.

19.1.5.5. TcpModbus Performance Determinants

The following description provides some insight into the possible sources of delay and error in an end-to-end TcpModbus exchange.

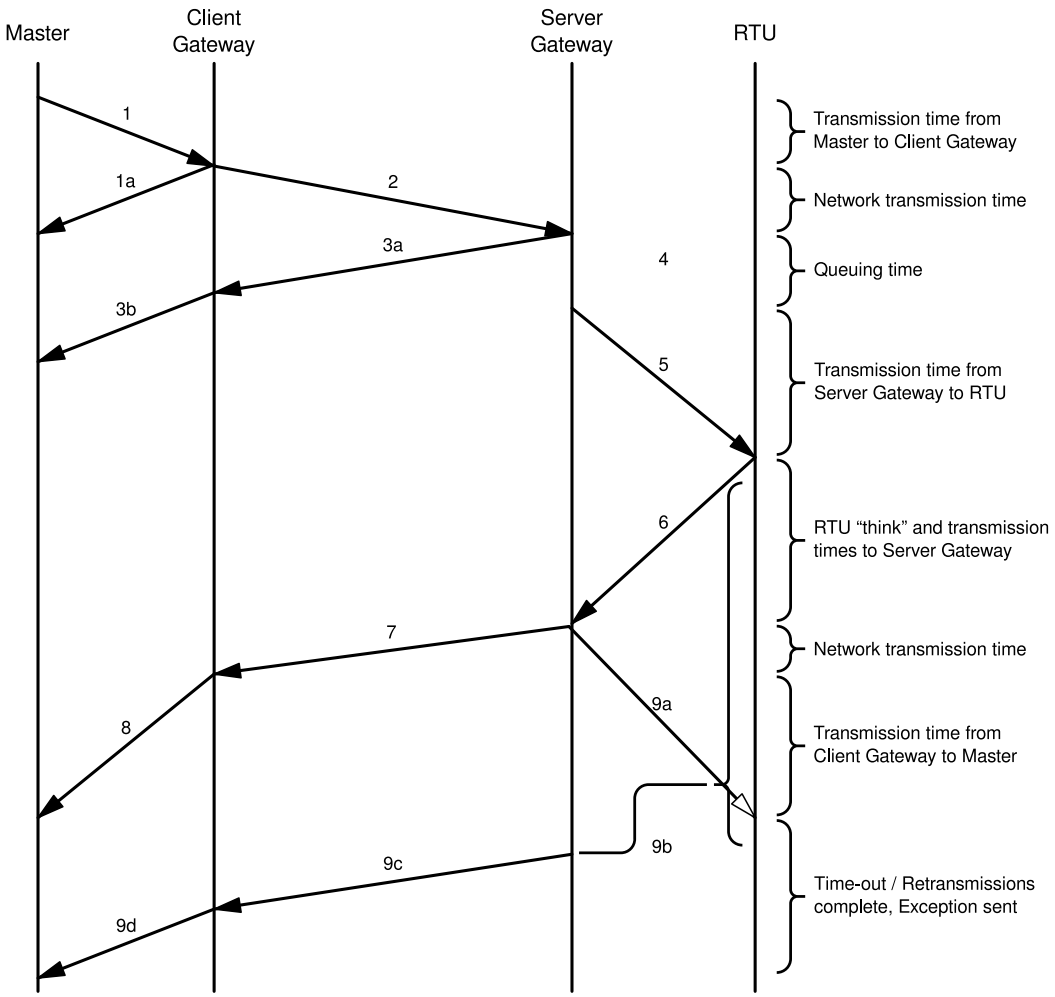

Figure 19.2. Sources of Delay and Error in an End to End Exchange

In step 1, the master issues a request to the Client Gateway. If the Client Gateway validates the message it will forward it to the network as step 2.

The Client Gateway can respond immediately in certain circumstances, as shown in step 1a. When the Client Gateway does not have a configuration for the specified RTU it will respond to the master with an exception using TcpModbus exception code 11 (“No Path”). When the Client Gateway has a configured RTU but the connection is not yet active it will respond to the master with an exception using TcpModbus exception code 10 (“No Response”). If the forwarding of TcpModbus exceptions is disabled, the client will not issue any responses.

Steps 3a and 3b represents the possibility that the Server Gateway does not have configuration for the specified RTU. The Server Gateway will always respond with a type 10 (“No Path”) in step 3a, which the client will forward in step 3b.

Step 4 represents the possibility of queuing delay. The Server Gateway may have to queue the request while it awaits the response to a previous request. The worst case occurs when a number of requests are queued for an RTU that has gone offline, especially when the server is programmed to retry the request upon failure.

Steps 5-8 represent the case where the request is responded to by the RTU and is forwarded successfully to the master. It includes the “think time” for the RTU to process the request and build the response.

Step 9a represents the possibility that the RTU is offline, the RTU receives the request in error or that the Server Gateway receives the RTU response in error. If the Server Gateway does not retry the request, it will issue an exception to the originator.

19.1.5.6. A Worked Example

A network is constructed with two Masters and 48 RTUs on four Server Gateways. Each of the Master is connected to a Client Gateway with a 115.2 Kbps line. The RTUs are restricted to 9600 bps lines. The network is Ethernet based and introduces an on average 3 ms of latency. Analysis of traces of the remote sites has determined that the min/max RTU think times were found to be 10/100 ms. What time-out should be used by the Master?

The maximum sized Modbus message is 256 bytes in length. This leads to a transmission time of about 25 ms at the Master and 250 ms at the RTU. Under ideal circumstances the maximum round trip time is given by: 25 ms (Master->client) + 3 ms (network delay) + 250 ms (server->RTU) + 100 ms (Think time) + 250 ms (RTU->server) + 3 ms (network delay) + 25 ms (client->Master). This delay totals about 650 ms.

Contrast this delay with that of a “quick” operation such as reading a single register. Both request and response are less than 10 bytes in length and complete (for this example) in 1 and 10 ms at the client and server. Assuming the RTU responds quickly, the total latency will approach 35 ms.

It is also necessary to take account such factors as the possibility of line errors and collisions between masters at the server.

The server may be configured to recover from a line error by retransmitting the request. Given a maximum frame transmission time of 250 ms and an RTU latency of 100 ms, it would be wise to budget 350 ms for each attempt to send to the RTU. Configuring a single retransmission would increase the end-to-end delay from about 650 ms to about 1000 ms.

The server can already be busy sending a request when the request of our example arrives. Using the figures from the above paragraph, the server being busy would increase the end-to-end delay from 1000 to 1350 ms.

The preceding analysis suggests that the Master should time-out at some time after 1350 ms from the start of transmission.

19.1.6. DNP (Distributed Network Protocol)

The RX1000 supports DNP 3.0, commonly used by utilities in process automation systems. DNP3 protocol messages specify source and destination addresses. A destination address specifies which device should process the data, and the source address specifies which device sent the message. Having both destination and source addresses satisfies at least one requirement for peer-to-peer communication since the receiver knows where to direct a response. Each device supporting the DNP protocol must have a unique address within the collection of devices sending and receiving DNP messages.

19.1.6.1. Address Learning for DNP

The RX1000 implements both local and remote address learning for DNP.

A local Device Address Table is populated with DNP Addresses learned for local and remote DNP devices. Each DNP address is associated with either a local serial port or a remote IP address.

When a message with an unknown DNP source address is received on a local serial port, the DNP source address and serial port number are entered into the Device Address Table. When a message with an unknown DNP source address is received from the IP network, on the IP interface that is configured as the DNP learning interface, the DNP source address and the IP address of the sender are entered into the Device Address Table.

When a message with an unknown DNP destination address is received on a local serial port, the message is sent in a UDP broadcast out the network interface configured as the DNP learning interface. When a message with an unknown DNP destination address is received from the IP network, it is sent to all local serial ports configured as DNP ports.

UDP transport is used during the DNP address learning phase.

All learned addresses will be kept in the Device Address Table, which is saved in non-volatile memory, which makes it unnecessary to repeat the DNP address learning process across a RX1000 reboot or accidental power loss.

An aging timer is maintained per DNP address in the table, and is reset whenever a DNP message is sent to or received for the specified address.

This learning facility makes it possible to configure the DNP3 protocol with a minimum number of parameters: a TCP/UDP port number, a learning network interface and an aging timer.

19.1.6.2. DNP Broadcast Messages

DNP addresses 65521 through 65535 are reserved as DNP3 broadcast addresses. The RX1000 supports DNP3 broadcast messages. DNP broadcast messages received on local serial ports are transmitted to all IP Addresses in the DNP Device Address Table (whether learned or statically configured). When a DNP broadcast message is received from the IP network, it is transmitted on all local serial ports configured as DNP ports.

19.2. Serial Protocol Configuration

Figure 19.3. Serial Protocols menu

To display the Serial Protocols menu, navigate to *interface/serial*.

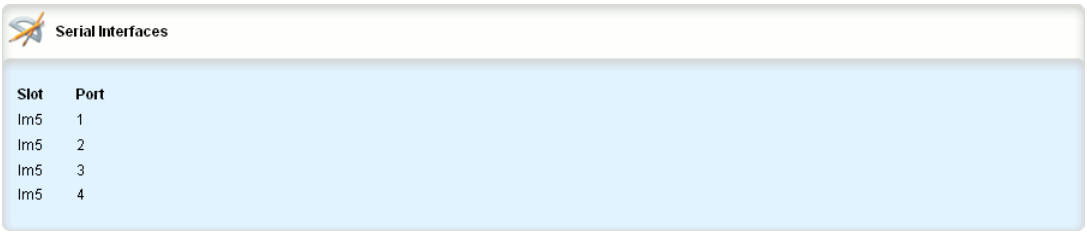

Figure 19.4. Serial Interfaces table

If data and ports have been configured, the Serial Interfaces table appears on the same screen as the *Serial Protocols* menu.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location provided on the silkscreen across the top of the device.

port

Synopsis: *integer*

The port number as seen on the front plate silkscreen of the switch (or a list of ports, if aggregated in a port trunk).

enabled

Synopsis: *boolean*

Default: *true*

Provides the option to enable or disable this interface. When unchecked(i.e disabled), the interface will prevent all frames from being sent and received on that interface.

alias

Synopsis: *A string*

The SNMP alias name of the interface

19.2.1. Assigning Protocols

Select the type of protocol to assign to a serial port.

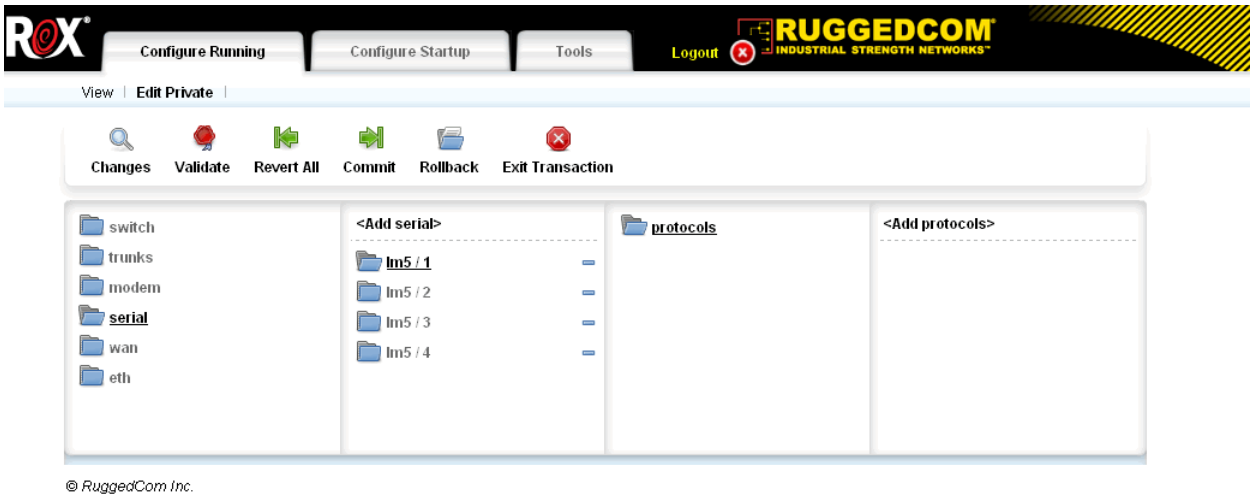

Figure 19.5. Adding a Protocol in the Edit Private screen

In Edit Private view, the <Add protocols> option can be clicked, which adds a protocol to a port.

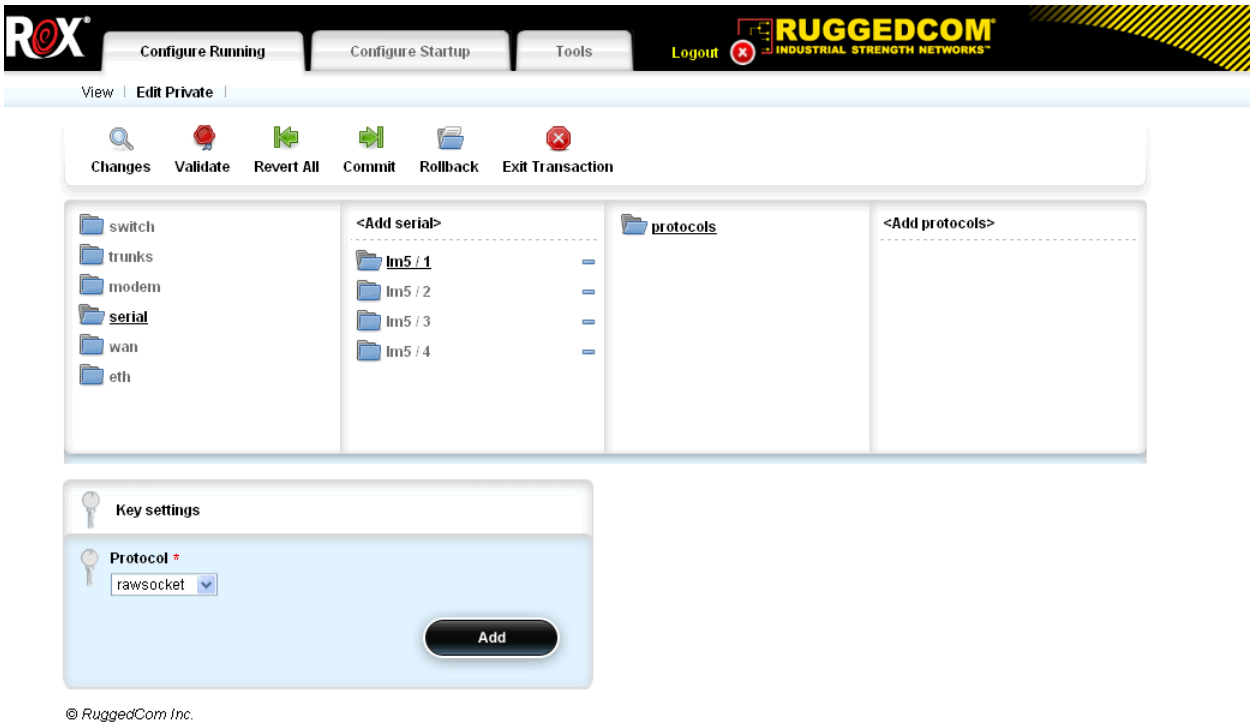

Figure 19.6. Selecting the Protocol Type in the Edit Private screen

Selecting a protocol type from the *Protocol* field in the Key Settings form associates a protocol with a serial port. Rawsocket, TcpModbus or DNP protocols can be selected. Unused ports should be left associated with "none". Changing an association immediately closes the calls of any protocols previously selected on the same port.

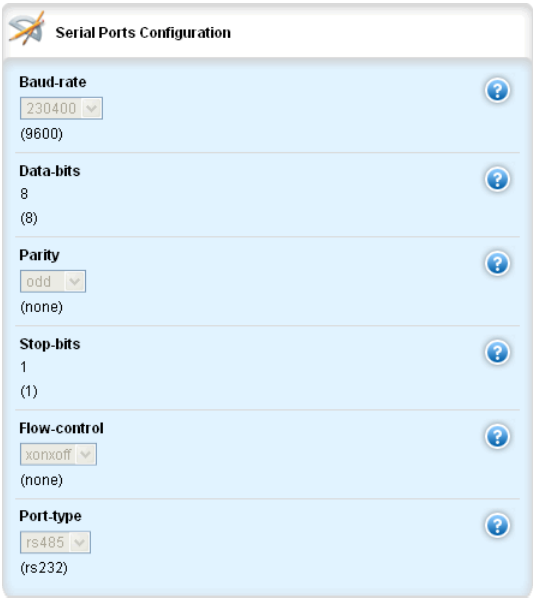

Figure 19.7. Serial Ports Configuration form

The Serial Interfaces form configures the serial settings and electrical protocol associated with a serial port. Changes are made immediately. To display this form, navigate to *interface/serial/{line module}*.

baud-rate

Synopsis: *string - one of the following keywords { 230400, 115200, 57600, 38400, 19200, 9600, 2400, 1200 }*

Default: *9600*

The baudrate selection of serial port

data-bits

Synopsis: *integer*

Default: *8*

The number of data bits

parity

Synopsis: *string - one of the following keywords { odd, even, none }*

Default: *none*

The parity of the serial port

stop-bits

Synopsis: *integer*

Default: *1*

The number of stop bits of the serial port

flow-control

Synopsis: *string - one of the following keywords { xonxoff, none }*

Default: *none*

Flow control of the serial port

port-type

Synopsis: *string - one of the following keywords { rs485, rs422, rs232 }*

Default: *rs232*

The type of serial port

Figure 19.8. Serial Protocols table

The Serial Protocols table displays the protocols configured. To display the Serial Interfaces table, navigate to *interface/serial/{line module}/protocols*.

protocol

Synopsis: *string - one of the following keywords { dnp, tcpmodbus, rawsocket }*

19.2.2. Setting Rawsockets

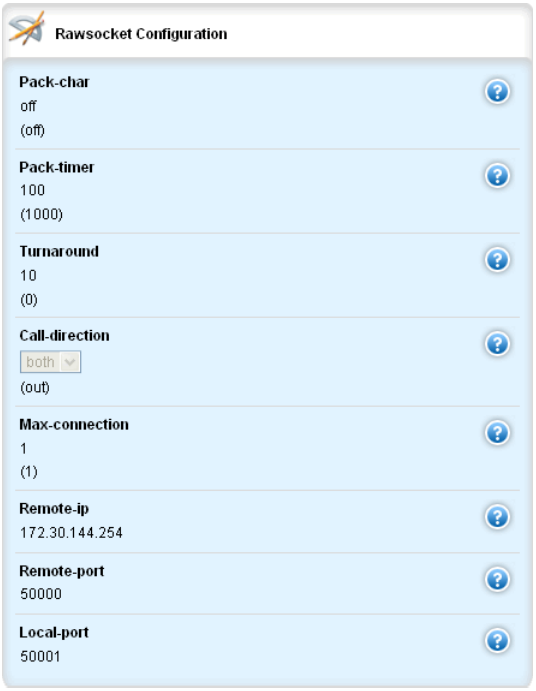

Figure 19.9. Rawsocket Configuration form

The Rawsocket Configuration form is used to configure the Raw Socket settings for each port. Changes are made immediately. To display the Rawsocket Configuration form, navigate to *interface/serial/{line module}/protocols/rawsocket/setrawsocket*.

pack-char

Synopsis: *string - the keyword { off }*

Synopsis: *integer*

Default: *off*

The numeric value of the ASCII character which will force forwarding of accumulated data to the network.

pack-timer

Synopsis: *integer*

Default: *1000*

The delay from the last received character until when data is forwarded

turnaround

Synopsis: *integer*

Default:

The amount of delay (if any) to insert between the transmissions of individual messages out the serial port

call-direction

Synopsis: *string - one of the following keywords { both, out, in }*

Default: *out*

Whether to accept an incoming connection, place an outgoing connection or do both

max-connection

Synopsis: *integer*

Default: *1*

The maximum number of incoming connections to permit when the call direction is incoming.

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address used when placing an outgoing connection.

remote-port

Synopsis: *integer*

The TCP destination port used in outgoing connections

local-port

Synopsis: *integer*

The local TCP port to use to accept incoming connections.

19.2.3. Setting TcpModbus

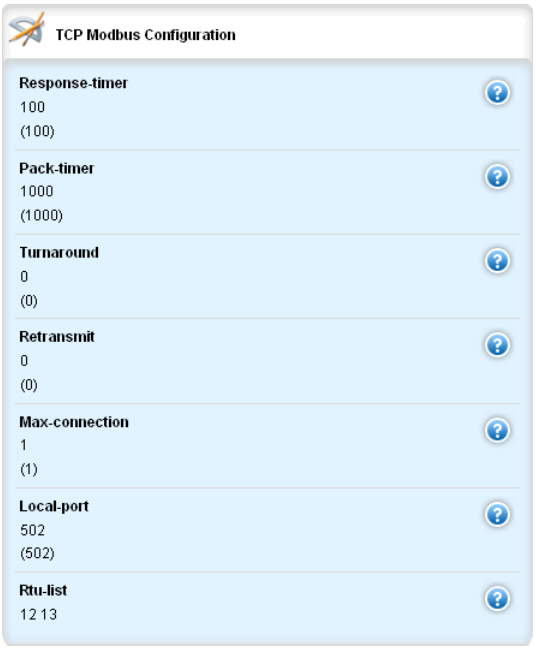

Figure 19.10. TCP Modbus Configuration form

The TCP Modbus Configuration form is used to configure the TcpModbus settings for each port. Changes are made immediately. To display the TCP Modbus Configuration form, navigate to *interface/serial/{line module}/protocols/tcpmodbus/settcpmodbus*.

response-timer

Synopsis: *integer*

Default: *100*

The maximum time from the last transmitted character of the outgoing poll until the first character of the response. If the RTU does not respond in this time, the poll will have been considered failed.

pack-timer

Synopsis: *integer*

Default: *1000*

The maximum allowable time to wait for a response to a Modbus

request to complete once it has started.

turnaround

Synopsis: *integer*

Default:

The amount of delay (if any) to insert after the transmissions of Modbus broadcast messages out the serial port.

retransmit

Synopsis: *integer*

Default:

The number of times to retransmit the request to the RTU before giving up.

max-connection

Synopsis: *integer*

Default: *1*

The maximum number of incoming connections.

local-port

Synopsis: *integer*

Default: *502*

The alternate local TCP port number. If this field is configured, a single connection (per serial port) may be made to this alternate port number. Note that TCP Modbus uses a default local port number of 502. There is no limit imposed on the number of connections to the default TCP port.

rtu-list

Synopsis: *string*

The ID of the RTU that is hooking up to the serial port.

The Modbus specification states the minimum time is about 640 character times at baud rates below 19200 Kbps and 256 char times + 192 ms at baud rates above 19200 Kbps. You may specify a larger value if you think your RTU will take longer to complete transmission than the calculated time.

19.2.4. Setting DNP

Figure 19.11. DNP Protocols Configuration form

The DNP Protocols Configuration form is used to configure the DNP settings for each port. To display the DNP Protocols Configuration form, navigate to *interface/serial/{line module}/protocols/dnp/setdnp*.

address-learning

Synopsis: *A string*

The interface to learn the RTU address from.

aging-timer

Synopsis: *integer*

Default: *1000*

The length of time which a learned DNP device in the Device Address Table may go without any DNP communication before it is removed from the table.

max-connection

Synopsis: *integer*

Default: *1*

The maximum number of incoming DNP connections.

Figure 19.12. DNP Device Address Table Configuration table

That Device Address table displays all currently known active DNP devices. To display the DNP Address Table Configuration form, navigate to *interface/serial/{line module}/protocols/dnp/setdnp/device-table*.

Figure 19.13. DNP Device Address Table Configuration form

To display the DNP Address Table Configuration form, navigate to *interface/serial/{line module}/protocols/dnp/setdnp/device-table/{number}*.

deviceAddress

Synopsis: *integer*

The local or remote DNP device address. The address may be that of a DNP device connected to a local serial port or one available via the serial port of a remote IP host.

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of the remote host that provides a connection to the DNP device with the configured address.Leave this field empty to forward DNP message that matches the configured address to local serial port

remote-device

Enable forwarding DNP message that matches the device address to remote-ip

19.3. Serial Protocol Statistics

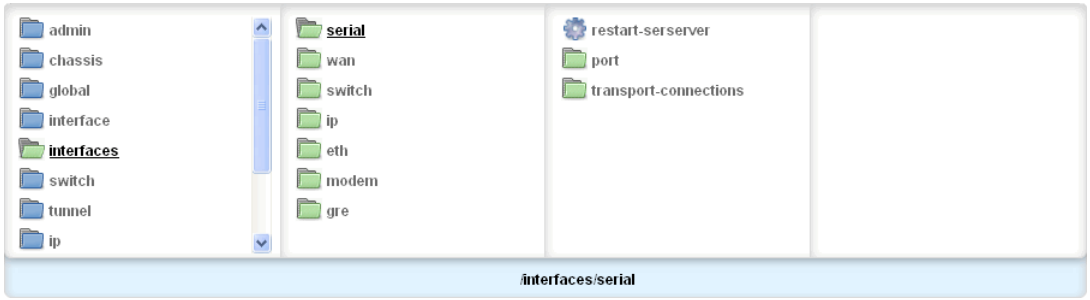

Figure 19.14. Serial Protocol Statistics menu

The Serial Protocol Statistics menu is accessible from the main menu under *interfaces/serial*.

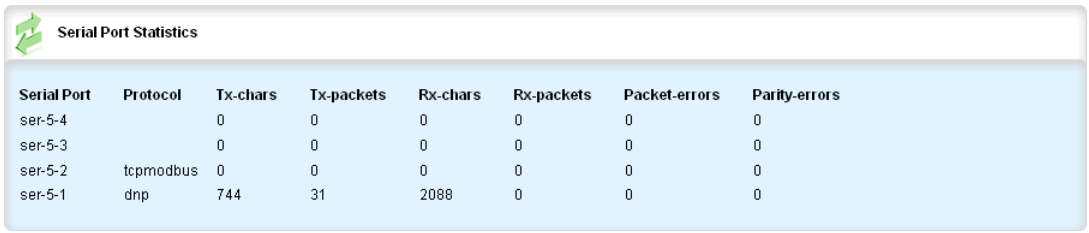

Serial Port	Protocol	Tx-chars	Tx-packets	Rx-chars	Rx-packets	Packet-errors	Parity-errors
ser-5-4		0	0	0	0	0	0
ser-5-3		0	0	0	0	0	0
ser-5-2	tcpmodbus	0	0	0	0	0	0
ser-5-1	dnp	744	31	2088	0	0	0

Figure 19.15. Serial Port Statistics table

To display the Serial Port Statistics table, navigate to *interfaces/serial/port*.

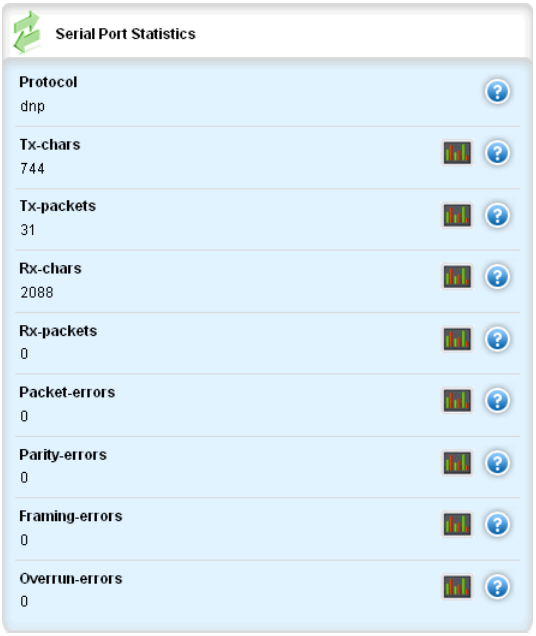

Figure 19.16. Serial Port Statistics form

To display the Serial Port Statistics form, navigate to *interfaces/serial/port* and then clicking on a linked submenu.

Serial Port

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

The serial interface name

media

Synopsis: *A string*

The type of port media { RS232 RS422 RS485 }. It provides the user with a description of the installed media type on the port for modular products. Please note that fiber media may be either Single Mode(SM), Multi Mode(MM), and may be Short Distance, Long Distance or Very Long Distance with connectors like LC, SC, ST, MTRJ etc. For the modules with SFP/GBICs, the media description is displayed per the SFF-8472 specification, if the transceiver is plugged into the module. E.g. 10/100/1000TX RJ45, 100FX SM SC, 10FX MM ST,1000SX SFP LC S SL M5.

speed

Synopsis: *string - one of the following keywords { 230.4K, 115.2K, 57.6K, 38.4K, 19.2K, 9.6K, 2.4K, 1.2K, 7.2M, 3.072M, 1.776M, 10G, 1G, 100M, 10M, 2.4M, 1.5M, auto }*

Speed (in Kilobits-per-second)

protocol

Synopsis: *A string*

The serial protocol assigned to this port

tx-chars

Synopsis: *unsigned integer*

The number of bytes transmitted over the serial port

tx-packets

Synopsis: *unsigned integer*

The number of packets transmitted over the serial port

rx-chars

Synopsis: *unsigned integer*

The number of bytes received by the serial port

rx-packets

Synopsis: *unsigned integer*

The number of packets received by the serial port

packet-errors

Synopsis: *unsigned integer*

The number of packet errors on this serial port

parity-errors

Synopsis: *unsigned integer*

The number of parity errors on this serial port

framing-errors

Synopsis: *unsigned integer*

The number of framing errors on this serial port

overrun-errors

Synopsis: *unsigned integer*

The number of overrun errors on this serial port

The Serial Port Statistics table and form present statistics of serial port activity and established connections. The Serial Port Statistics table provides a record for each active serial port. The number of historical received and transmitted characters as well as errors will be displayed.

19.3.1. Transport Connections

Transport Connection Statistics

Name	Remote-ip	Remote TCP/UDP port	Local TCP/UDP port	Transport	Rx-packets	Tx-packets	Status
ser-5-1/1	10.200.21.143	41865	20000	TCP	1	0	Active

Figure 19.17. Transport Connections Statistics table

The Transport Connection Statistics table reflects established TCP connections. To display the Transport Connections Statistics table, navigate to *interfaces/serial/transport-connections*.

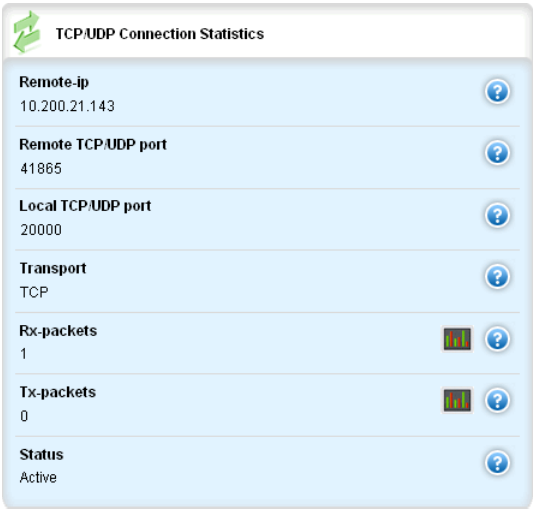

Figure 19.18. TCP/UDP Connection Statistics form

index

Synopsis: *A string*

The transport connection index

remote-ip

Synopsis: *A string*

The IP address of the remote serial server

Remote TCP/UDP port

Synopsis: *integer*

The port of the remote serial server

Local TCP/UDP port

Synopsis: *integer*

The local port for the incoming connection

transport

Synopsis: *A string*

The transport protocol (UDP or TCP) for this serial port

rx-packets

Synopsis: *unsigned integer*

The number of packets received from TCP/UDP

tx-packets

Synopsis: *unsigned integer*

The number of packets transmitted to TCP/UDP

target-port

Synopsis: *A string*

Target serial port

status

Synopsis: *A string*

The connection status of the serial port

19.4. Restarting the Serial Server

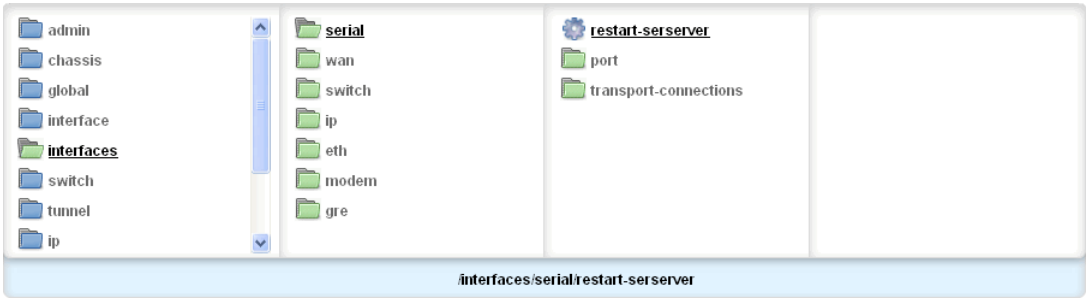

Figure 19.19. Restart Serserver menu

The path to the Restart Serserver menu is *interfaces/serial/restart-serserver*. To restart the serserver, click on the *restart-serserver* trigger action and the click the *Perform* button on the Trigger Action form.

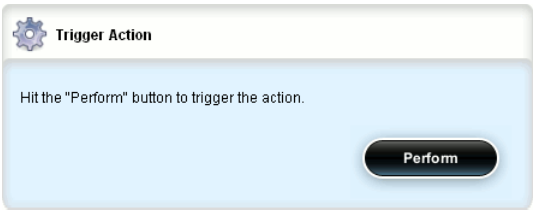

Figure 19.20. Restart Serserver Trigger Action

19.5. Resetting Ports

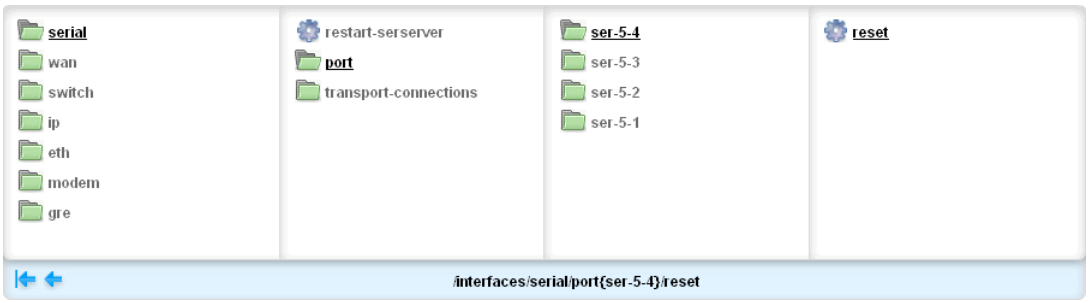

Figure 19.21. Reset Ports menu

The path to the Reset Ports menu is *interfaces/serial/port/{line module}/reset*. To reset the ports, click on the *reset* trigger action and then click the *Perform* button on the Reset Trigger Action form.

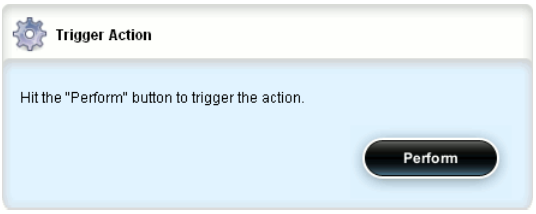

Figure 19.22. Reset Ports Trigger Action

20. WAN

20.1. T1/E1 Fundamentals

A T1 line is a communications circuit using the Digital Signal 1 (DS1) signalling scheme. DS1 allows 24 “timeslots” of 64 Kbps DS0 information, along with 8 Kbps of signalling information, to be multiplexed onto a 1544 Kbps circuit.

The 24 DS0s can be used individually as standalone channels, bonded into groups of channels, or can be bonded to form a single 1536 Kbps channel, referred to as a clear channel. Not all channels need be used. It is quite common to purchase a number of channels of 64Kbps bandwidth and leave the remainder unused; this is known as fractional T1.

The telephone network terminates the T1 line and maps each of the channels through the T1 network to a chosen T1 line. Individual and bonded DS0s from more than one remote T1 can be aggregated into a full T1 line. This is referred to as central site concentration.

The T1 line itself is referred to as the *physical interface*. Groups of DS0s form channels and the protocols that run on the channels are known as *logical interfaces*. The RuggedRouter® provides the ability to operate Frame Relay or PPP over logical interfaces.

An E1 line is a communications circuit conforming to European standards with 32 64 Kbps channels, one of which is usually reserved for signalling information.

20.1.1. Frame Relay

Frame Relay is a packet switching protocol for use over the WAN. The RuggedRouter® provides the ability to construct point-to-point IP network connections over Frame Relay.

Each Frame Relay interface provides a link between a local and a peer station. One of the stations must be configured as a Data Communications Equipment (DCE) device, often referred to as a Switch. The the peer station must be configured as a Data Terminal Equipment (DTE) device, often referred to as Customer Premises Equipment (CPE). The DCE is responsible for managing the link, advertising connections to the DTE, and switching packets between connections. The DTE raises individual connections and sends data on them. A frame relay switch is usually configured as DCE, and the RX1000 is configured as DTE.

When using a T1/E1 line to access a public Frame Relay provider, configure the router as a DTE.

Unlike PPP, a Frame Relay link can provide multiple connections. Each connection is identified by a Data Link Connection Identifier (DLCI) and must match at the DCE and DTE. The use of multiple connections can support meshed network interconnections and disaster recovery.

20.2. T3/E3 Fundamentals

A T3 line is communications circuit using a Digital Signal 3 (DS3) signalling scheme. DS3 allows 672 “timeslots” of 64 Kbps DS0 information to be multiplexed onto a 44.736 Mbps circuit.

E3 refers to the ITU standard corresponding to the mainly North American T3 standard. E3 calls for 512 DS0-equivalent timeslots multiplexed onto a 34.368 Mbps circuit.

RuggedRouter® provides the ability to operate Frame Relay or PPP over your physical T3/E3 interfaces.

Channel groups and fractional lines are not supported on RuggedRouter® T3 and E3 interfaces.

20.3. WAN Configuration

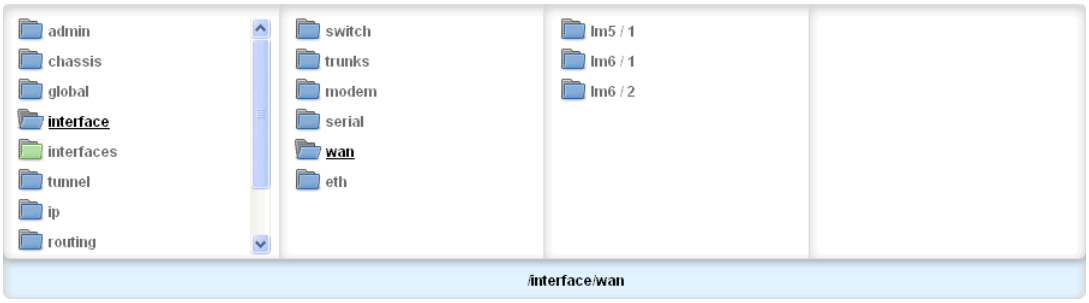

Figure 20.1. WAN menu

To display the WAN menu, navigate to *interface/wan*. The WAN Slot Port Settings table appears on the same page as the WAN menu.

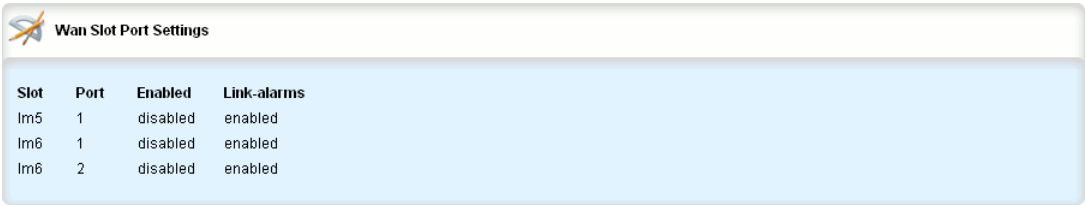
Wan Slot Port Settings
Slot Port Enabled Link-alarms
lm5 1 disabled enabled
lm6 1 disabled enabled
lm6 2 disabled enabled

Figure 20.2. Interface WAN Slot Port Settings table

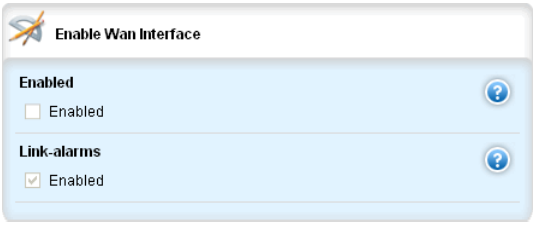

Figure 20.3. Enable WAN Interface form

The path to the Enable WAN Interface form is *interface/wan/{line module}*.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

The name of the module location for the WAN card.

port

Synopsis: *integer*

The port number on the WAN card.

enabled

Synopsis: *boolean*

Default: *false*

Enables this WAN port.

link-alarms

Synopsis: *boolean*

Default: *true*

Disabling link-alarms will prevent alarms and LinkUp and LinkDown SNMP traps from being sent for that interface. Link alarms may also be controlled for the whole system under admin / alarm-cfg.

alias

Synopsis: *A string*

The SNMP alias name of the interface

20.3.1. T1 Parameters

You can configure T1 parameters for a WAN port. The path to the T1 Parameters form is *interface/wan/{line module}/T1*.

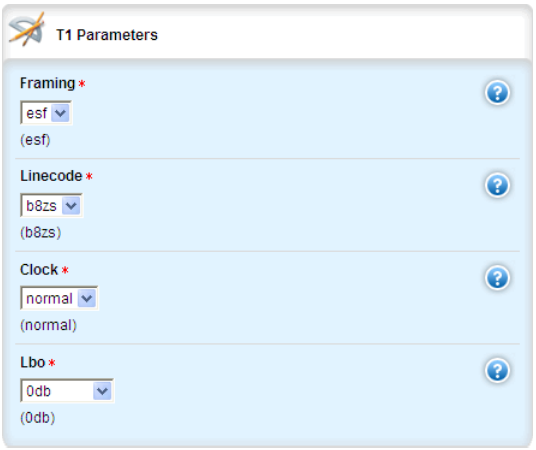

Figure 20.4. T1 Parameters form

framing

Synopsis: *string - the keyword { esf }*

Default: *esf*

The frame format.

linecode

Synopsis: *string - the keyword { b8zs }*

Default: *b8zs*

The line encoding/decoding scheme.

clock

Synopsis: *string - one of the following keywords { master, normal }*

Default: *normal*

Serial clocking mode: master or normal.

- master : provide serial clock signal.
- normal : accept external clock signal.

lbo

Synopsis: *string - one of the following keywords { 550-660ft, 440-550ft, 330-440ft, 220-330ft, 110-220ft, 0-110ft, 22.5db, 15db, 7.5db, 0db }*

Default: *0db*

Line Build Out: tunes the shape of the T1 pulses and adjusts their amplitude depending upon distances and the desired attenuation.

20.3.2. E1 Parameters

You can configure E1 Parameters for a WAN port. The path to the E1 Parameters form is *interface/wan/{line module}/E1*.

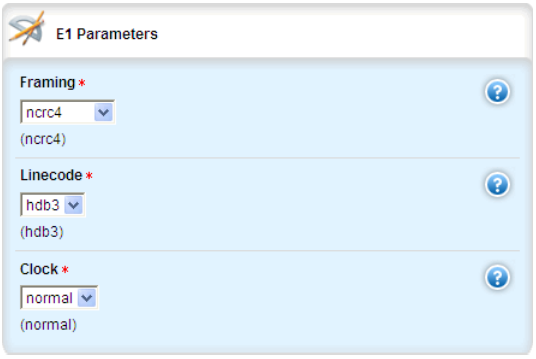

Figure 20.5. E1 Parameters form

framing

Synopsis: *string - one of the following keywords { unframed, crc4, ncrc4 }*

Default: *ncrc4*

The frame format.

linecode

Synopsis: *string - the keyword { hdb3 }*

Default: *hdb3*

A line encoding/decoding scheme.

clock

Synopsis: *string - one of the following keywords { master, normal }*

Default: *normal*

Serial clocking mode: master or normal.

- master : provide serial clock signal.
- normal : accept external clock signal.

20.3.3. T3/E3 Parameters

You can configure T3 parameters for a WAN port. The path to the T3 Parameters form is *interface/wan/{line module}/t3*.

Figure 20.6. T3 Parameter form

framing

Synopsis: *string - one of the following keywords { m13, c-bit }*

Default: *c-bit*

The frame format.

linecode

Synopsis: *string - one of the following keywords { ami, b3zs }*

Default: *b3zs*

A line encoding/decoding scheme.

clock

Synopsis: *string - one of the following keywords { master, normal }*

Default: *normal*

Serial clocking mode: master or normal.

- master : provide serial clock signal.
- normal : accept external clock signal.

20.3.4. E3 Parameters

You can configure E3 parameters for a WAN port. The path to the E3 Parameters form is *interface/wan/{line module}/e3*.

Figure 20.7. E3 Parameter form

framing

Synopsis: *string - one of the following keywords { g.832, g.751 }*

Default: *g.751*

The frame format.

linecode

Synopsis: *string - one of the following keywords { ami, hdb3 }*

Default: *hdb3*

Line encoding/decoding scheme.

clock

Synopsis: *string - one of the following keywords { master, normal }*

Default: *normal*

Serial clocking mode: master or normal.

- master : provide serial clock signal.
- normal : accept external clock signal.

20.3.5. Configuring Protocols

The path to the T1 Channels and Associated Time Slots table is */interface/wan{line module and port}/t1/channel*.

Figure 20.8. T1 Channels and Associated Time Slots table

The path to the T1 Time Slots form is */interface/wan{line module and port}/t1/channel{channel id}*.

20.3.5.1. Adding Channels

You can configure one or more channels over one t1/e1 physical interface, and each channel can have one or more time slots.

> *A maximum of 24 timeslots (all of the timeslots) can be allocated to a channel. However, the same time slots cannot be assigned to two channels belonging to the same physical interface.*

To add a channel, follow these steps:

1. Enter edit mode and navigate to */interface/wan{lm6 2}/t1* or */interface/wan{lm6 2}/e1*.
2. Click the icon beside *t1* or *e1*.
3. Click channel and <Add channel>. The Key settings form appears.
4. In the Key settings form, enter a number in the range of 1 to 24 and click **Add**. The T1 Time Slots form appears.

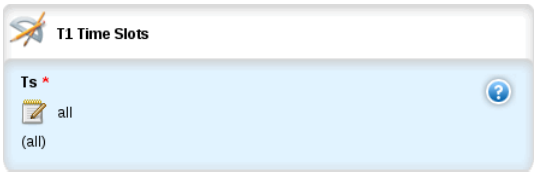

Figure 20.9. T1 Time Slots form

ts

Synopsis: *A string conforming to: "(all|[0-9,.]+)"*

Default: *all*

Time slots for this channel. Format: the string 'all', or a comma-separated list of numbers in the range of 1 to 24.

To specify a range of numbers, separate the start and end of the range with '..'

Example 1: 1,2,3 and 1..3 both represent time slots 1 through 3.

Example 2: 1,2,5..10,11 represents time slots 1, 2, 5, 6, 7, 8, 9, 10, and 11.

5. Commit the changes.

20.3.5.2. Adding Connections

After adding a channel, add connections to the channel. To add connections, enter edit mode and navigate to *interface/wan/{line module}/t1 (or e1)/{line module}/connection*.

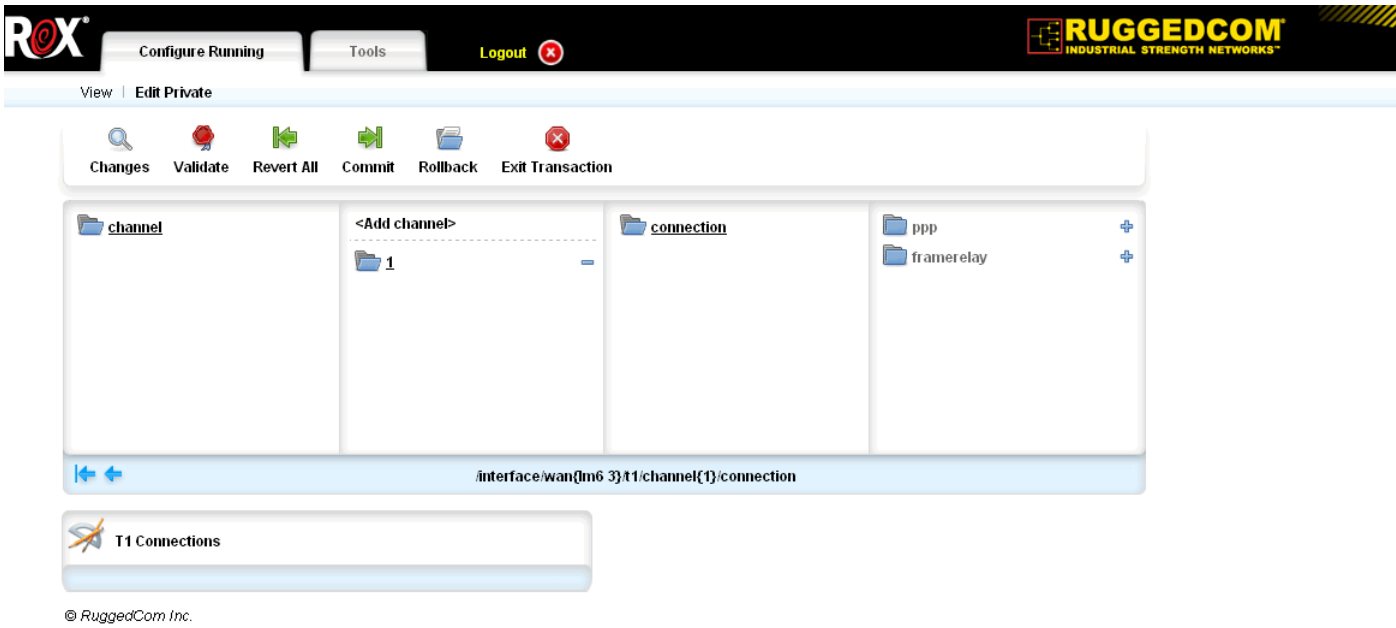

Figure 20.10. Adding a Connection

20.3.5.3. Configuring Frame Relay

From the *connection* submenu (see Figure 20.10, "Adding a Connection"), add a framerelay connection by clicking on the plus sign icon next to the *framerelay* submenu. Configure the parameters in the Frame Relay Parameter form.

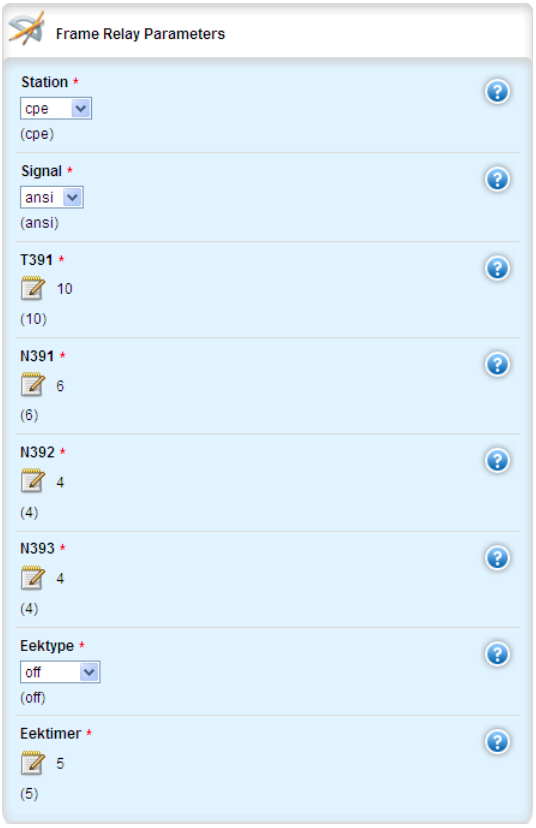

Figure 20.11. Frame Relay Parameter form

station

Synopsis: *string - one of the following keywords { switch, cpe }*

Default: *cpe*

The behavior of the frame relay connection, i.e. CPE (Customer Premises Equipment) or as a switch.

signal

Synopsis: *string - one of the following keywords { none, q933, lmi, ansi }*

Default: *ansi*

The frame relay link management protocol used.

t391

Synopsis: *integer*

Default: *10*

(Link Integrity Verification polling) Indicates the number of seconds between transmission of in-channel signaling messages. Valid for cpe.

t392

Synopsis: *integer*

Default: *16*

(Verification of polling cycle) Indicates the expected number of seconds between reception of in-channel signaling messages transmitted by cpe. Valid for Switch.

n391

Synopsis: *integer*

Default: *6*

Defines the frequency of transmission of full status enquiry messages. Valid for CPE.

n392

Synopsis: *integer*

Default: *4*

The number of error events (enumerated by n393) for which the channel is declared inactive; valid for either cpe or Switch.

n393

Synopsis: *integer*

Default: *4*

The number of error events on the frame relay channel; valid for either

cpe or switch.

eektype

Synopsis: *string - one of the following keywords { request, off }*

Default: *off*

Enables or disables EEK function.

eektimer

Synopsis: *integer*

Default: *5*

The number of seconds to wait before the next EEK message is sent.

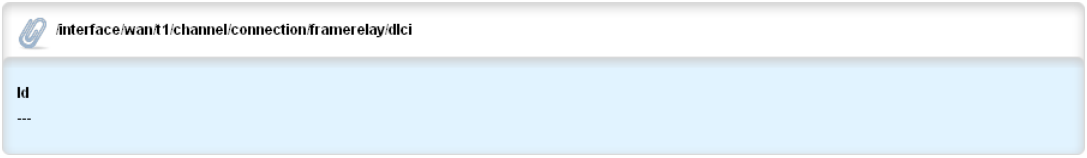

Figure 20.12. Connection Frame Relay DLCI table

This table displays the data link connection.

id

Synopsis: *integer*

DLCI (data link connection identifier) Number.

on-demand

This interface is up or down on demand of link fail over.

20.3.5.4. Configuring PPP

From the *connection* submenu (see Figure 20.10, “Adding a Connection”), add a PPP connection by clicking on the plus sign icon next to the *PPP* submenu. Click the *Commit* button and then click the *Exit Transaction* button. A PPP connection has now been added.

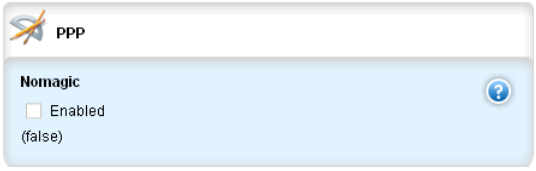

Figure 20.13. PPP form

nomagic

Synopsis: *boolean*

Default: *false*

Disables the Magic Number. (Valid on RX1000 only)

on-demand

This interface is up or down on demand of link fail over.

20.3.5.5. Configuring MLPPP

The PPP Multilink Protocol, also known as Multilink PPP or MLPPP, is defined in Internet RFC 1990. Its purpose is to combine two or more PPP links into a single “bundle” to provide more bandwidth for a point-to-point connection.

PPP Multilink must be supported on both sides of the link and may be used if there is more than one PPP link connecting the two endpoints. PPP works by multiplexing data on a per-packet basis to transmit across multiple PPP links. PPP uses sequence numbering to attempt to preserve the order of packets transmitted across the bundle.

ROX™ is capable of running PPP Multilink over two or more T1/E1 links, but is capable of defining only one MLPPP bundle.

For optimal PPP Multilink operation, ensure that each link in the MLPPP bundle has the same bandwidth: the number of time slots, the clocking mode, and the rate for each T1/E1 link used by PPP Multilink should be the same.

MLPPP interfaces use the following naming convention: **{physical interface}-mlppp-{bundle number}**. For example, **te1-mlppp-01** is an MLPPP interface on physical interface **te1** with bundle number **01**.

MLPPP interfaces use the following naming convention: **{physical interface}-mlppp-{bundle number}**. For example, **te1-mlppp-01** is an MLPPP interface on physical interface **te1** with bundle number **01**.

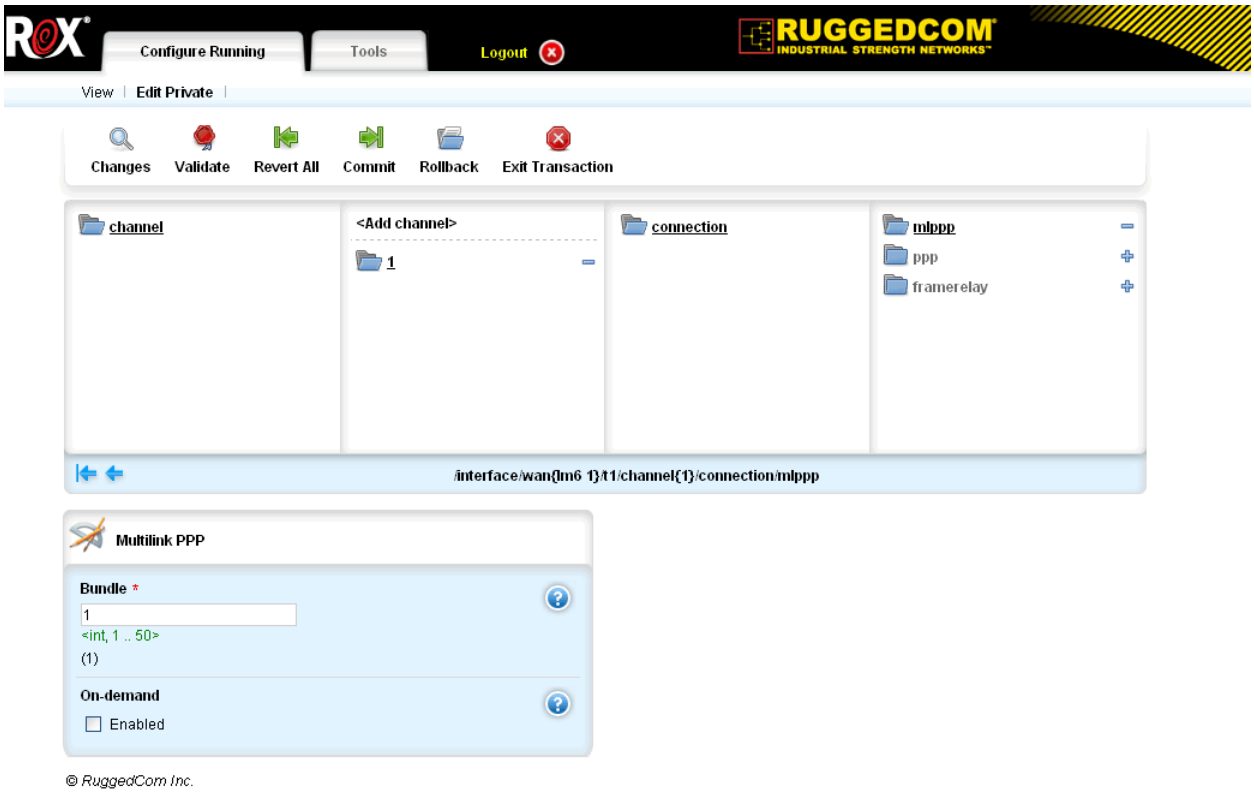

Figure 20.14. Adding an MLPPP Connection

From the *connection* submenu (see Figure 20.10, “Adding a Connection”), add an MLPPP connection by clicking the icon beside the *MLPPP* submenu. The Multilink PPP form appears.

Multilink PPP Form Parameters:

bundle

Synopsis: *integer*

Default: *1*

The bundle number

on-demand

This interface is up or down on demand of link fail over.

In the *Bundle* field, enter a bundle number (the default is 1). An MLPPP connection is added and an interface for this connection appears under the *ip* menu.

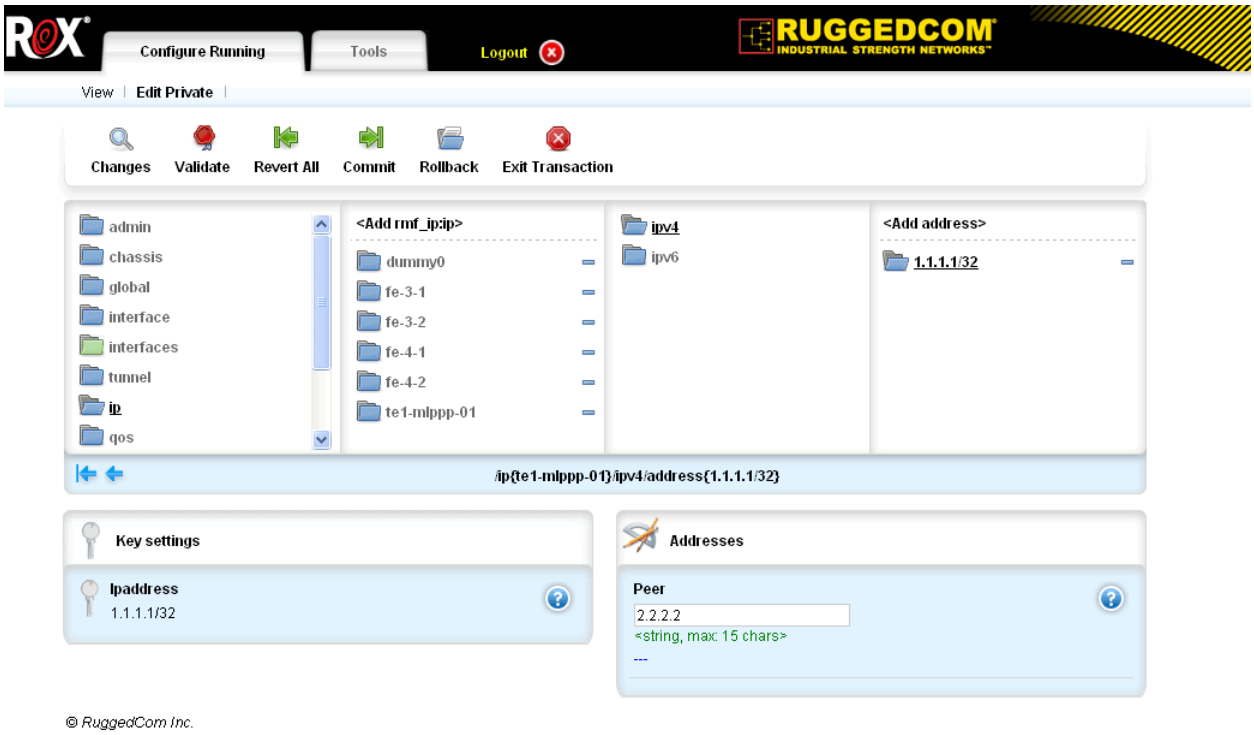

Figure 20.15. Adding IP and Remote Addresses

Navigate to *ip/{mlppp-interface}/ipv4* and click *Add address*. The Key settings form appears. In the *IPaddress* field, enter the IP address and click **Add**. The Addresses form appears.

In the *Peer* field, enter the remote IP address.

Click the **Commit** button and then click the **Exit Transaction** button.

> *You can add multiple PPP interfaces to a MLPPP link by configuring the same bundle number across all t1/e1 channels that are part of MLPPP.*

20.3.5.6. Naming and IP Address Assignment for Logical Interfaces

All interface identifiers use the following naming convention:

te*N*-*[physical interface number]*-c*[channel identifier]*[*connection number*]

- **te*N*** identifies the interface as a WAN interface (for example, te1 represents t1/e1).
- ***[physical interface number]*** identifies the physical interface.
- **c*[channel number]*** identifies the channel number.
- ***[connection identifier]*** identifies the type of connection. The connection identifier can be any of the following:
  - **ppp**
  - **mlppp** with a bundle number
  - **f*N*** for frame relay, where ***N*** is the Data Link Connection Identifier (DLCI), which is a four-digit number in the range of 0016 to 1007.

Examples:

- **te1-4-1c03ppp** represents t1/e1 in slot 4 port 1, channel 3 is configured for ppp

- **te1-4-1c04f0101** represents t1/e1 in slot 4 port 1, channel 4 is configured for FR DLCI 101
- **te1-mlppp-01** represents mlppp bundle 1

You can assign an IP Address for a t1/e1 logical interface in the IP submenu available from the main menu.

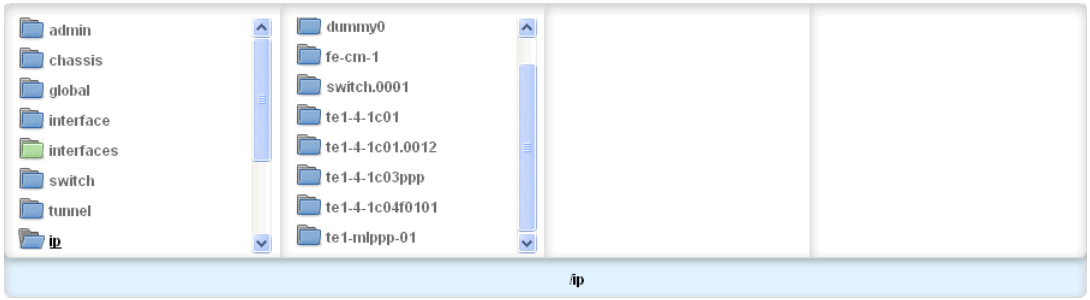

Figure 20.16. T1/E1 Interfaces under the IP submenu

All logical interfaces over t1/e1 are considered to be point to point links. Therefore the only acceptable subnet mask for the point-to-point link is /32.

20.3.6. Loopback Test

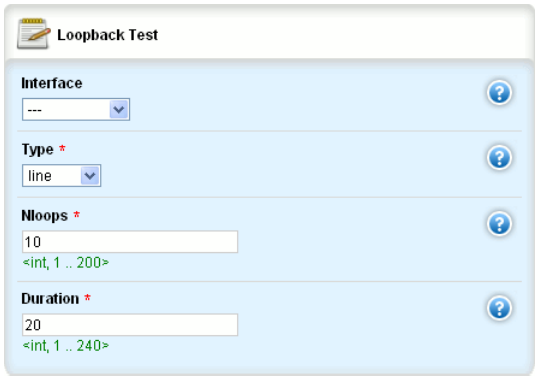

Figure 20.17. Loopback Test Forms

The path to the Loopback Test forms is *interfaces/wan/loopback*. In the Loopback Test form, select the *Interface* and *Type* and then set the *Nloops* and *Duration* parameters. To start a loopback test, click the **Perform** button on the trigger action form.

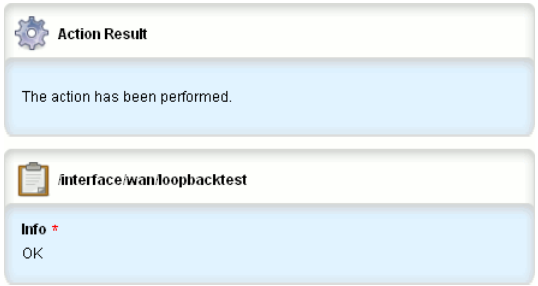

Figure 20.18. Loopbacktest Results

After launching the Loopback test, the Action Result form and the Loopbacktest form appear to confirm that the test has been performed and whether it was successful or not.

20.4. Statistics

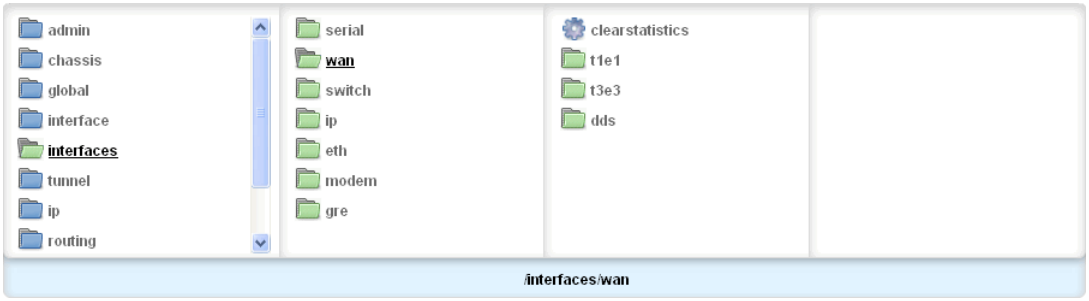

Figure 20.19. WAN Statistics menu

The WAN statistics are accessible via the *WAN Statistics* menu. The path to this menu is *interfaces/wan*.

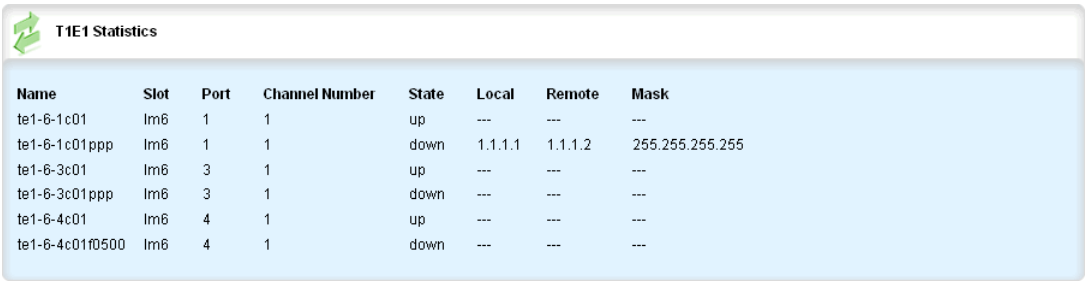
T1E1 Statistics
Name Slot Port Channel Number State Local Remote Mask
te1-6-1c01 lm6 1 1 up --- --- ---
te1-6-1c01ppp lm6 1 1 down 1.1.1.1 1.1.1.2 255.255.255.255
te1-6-3c01 lm6 3 1 up --- --- ---
te1-6-3c01ppp lm6 3 1 down --- --- ---
te1-6-4c01 lm6 4 1 up --- --- ---
te1-6-4c01f0500 lm6 4 1 down --- --- ---

Figure 20.20. T1E1 Statistics table

The path to this table is *interfaces/wan/t1e1/{line module}/t1e1*.

20.4.1. Physical Layer-related Statistics

Figure 20.21. Receiving Errors Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/receive-error*.

Over Run

Synopsis: *unsigned integer*

The number of receiver overrun errors.

CRC Error

Synopsis: *unsigned integer*

The number of receiver CRC errors.

Abort

Synopsis: *unsigned integer*

The number of receiver abort errors.

Corruption

Synopsis: *unsigned integer*

The number of receiver corruption errors.

PCI Error

Synopsis: *unsigned integer*

The number of receiver PCI errors.

DMA Error

Synopsis: *unsigned integer*

The number of receiver DMA descriptor errors.

Length Error

Synopsis: *unsigned integer*

Length errors.

Frame Error

Synopsis: *unsigned integer*

Frame errors.

Figure 20.22. T1E1 Receiving Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/receive-stats*.

Frames

Synopsis: *unsigned integer*

The number of frames received.

Bytes

Synopsis: *unsigned integer*

The number of bytes received.

Frames Discarded as Link Inactive

Synopsis: *unsigned integer*

Received frames that were discarded (link inactive).

Figure 20.23. T1E1 Receiving Statistics Form 2

The path to this form is *interfaces/wan/t1e1/{line module}/receive*.

Bytes

Synopsis: *unsigned long integer*

Number of bytes received.

Packets

Synopsis: *unsigned long integer*

Number of packets received.

Errors

Synopsis: *unsigned integer*

Number of error packets received.

Dropped

Synopsis: *unsigned integer*

Number of packets dropped.

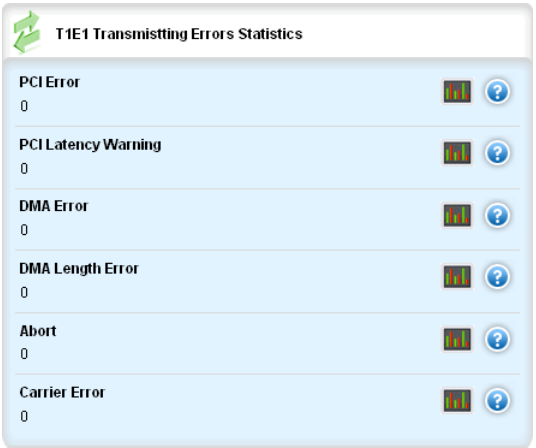

Figure 20.24. T1E1 Transmitting Errors Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/transmit-error*.

PCI Error

Synopsis: *unsigned integer*

The number of transmitter PCI errors.

PCI Latency Warning

Synopsis: *unsigned integer*

The number of transmitter PCI latency warnings.

DMA Error

Synopsis: *unsigned integer*

The number of transmitter DMA descriptor errors.

DMA Length Error

Synopsis: *unsigned integer*

The number of transmitter DMA descriptor length errors.

Abort

Synopsis: *unsigned integer*

Abort errors.

Carrier Error

Synopsis: *unsigned integer*

Carrier errors.

Figure 20.25. T1E1 Transmitting Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/transmit-stats*.

Frames

Synopsis: *unsigned integer*

The number of frames transmitted.

Bytes

Synopsis: *unsigned integer*

The number of bytes transmitted.

Realigned

Synopsis: *unsigned integer*

Transmits frames that were realigned.

Figure 20.26. T1E1 Transmitting Statistics Form 2

The path to this form is *interfaces/wan/t1e1/{line module}/transmit*.

Bytes

Synopsis: *unsigned long integer*

Number of bytes transmitted.

Packets

Synopsis: *unsigned long integer*

Number of packets transmitted.

Errors

Synopsis: *unsigned integer*

Number of error packets.

Dropped

Synopsis: *unsigned integer*

Number of packets dropped.

Collision

Synopsis: *unsigned integer*

Number of collisions detected.

Figure 20.27. T3 and E3 Alarm form

Alarm physical statistics are displayed in the T3 and E3 Alarm forms. The path to these forms is *interfaces/wan/t3e3/{line module}/alarm*.

ais

Synopsis: *string*

AIS (Alarm Indication Signal) alarm.

los

Synopsis: *string*

LOS (Loss Of Signal) alarm.

oof

Synopsis: *string*

OOF (Out Of Frame) alarm.

yel

Synopsis: *string*

YELLOW (RAI [Remote Alarm Indication] failure) alarm.

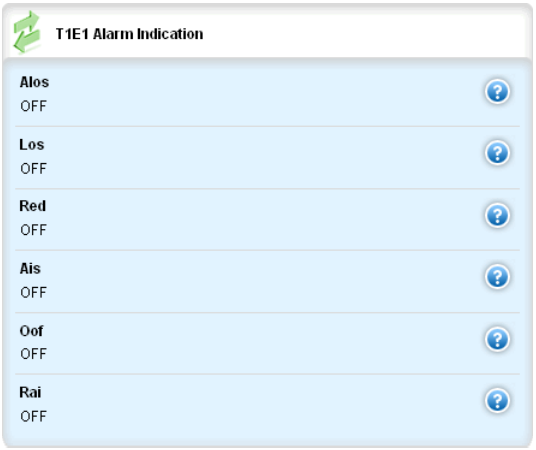

Figure 20.28. T1E1 Alarm Indication form

Alarm physical statistics are displayed in the T1E1 Alarm Indication form. The path to this form is *interfaces/wan/t1e1/{line module}/alarm*.

alos

Synopsis: *string*

ALOS (Loss of Signal) alarm.

los

Synopsis: *string*

LOS (Loss Of Signal) alarm.

red

Synopsis: *string*

RED (red alarm is a combination of a LOS or an OOF failure) alarm.

ais

Synopsis: *string*

AIS (Alarm Indication Signal) alarm.

oof

Synopsis: *string*

OOF (Out Of Frame) alarm.

rai

Synopsis: *string*

RAI (Remote Alarm Indication) alarm.

20.4.2. Protocol-related Statistics

20.4.2.1. PPP Statistics Summary

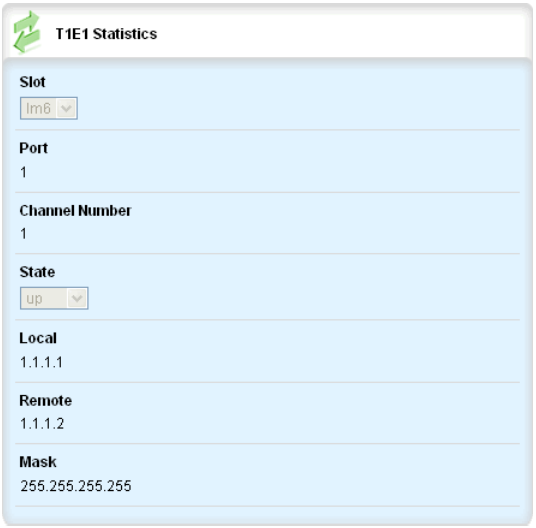

Figure 20.29. T1E1 Statistics form

The T1E1 Statistics form displays PPP statistics and physical statistics. The path to this form is *interfaces/wan/t1e1/{line module}*.

Figure 20.30. PPP Receiving Protocol Statistics form

The PPP Receiving Protocol Statistics form displays PPP receiving statistics. The path to this form is *interfaces/wan/t1e1/{line module}/ppp-stats*.

LCP

Synopsis: *unsigned integer*

The number of LCP (Link Control Protocol) packets.

PAP

Synopsis: *unsigned integer*

The number of PAP (Password Authentication Protocol) packets.

CHAP

Synopsis: *unsigned integer*

The number of CHAP (Challenge Handshake Authentication Protocol) packets.

IPCP

Synopsis: *unsigned integer*

The number of IPCP (Internet Protocol Control Protocol) packets.

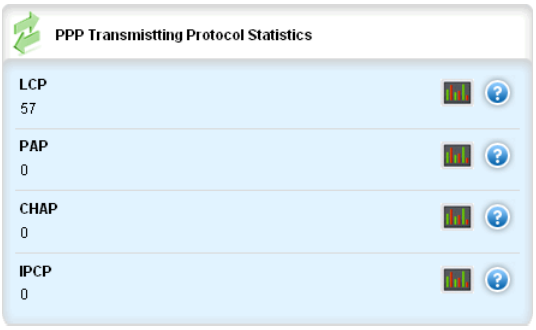

Figure 20.31. PPP Transmitting Protocol Statistics form

The PPP Receiving Protocol Statistics form displays PPP transmitting statistics. The path to this form is *interfaces/wan/t1e1/{line module}/ppp-stats*. Additional statistics forms can be found under this *ppp-stats* submenu.

LCP

Synopsis: *unsigned integer*

The number of LCP (Link Control Protocol) packets.

PAP

Synopsis: *unsigned integer*

The number of PAP (Password Authentication Protocol) packets.

CHAP

Synopsis: *unsigned integer*

The number of CHAP (Challenge Handshake Authentication Protocol) packets.

IPCP

Synopsis: *unsigned integer*

The number of IPCP (Internet Protocol Control Protocol) packets.

20.4.2.2. MLPPP Statistics

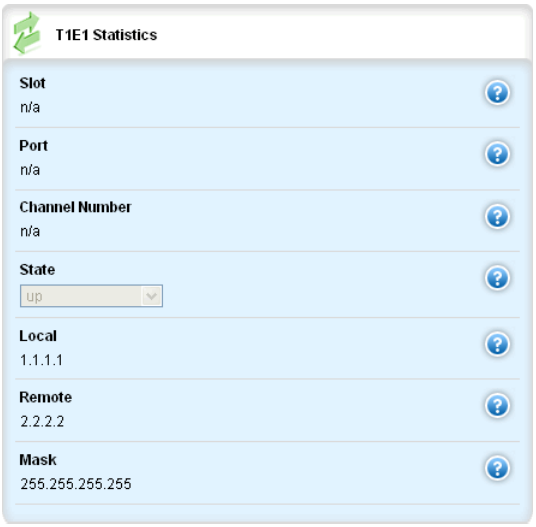

Figure 20.32. T1E1 Statistics form

The path to this form is *interfaces/wan/t1e1/{mlppp line module}*.

name

Synopsis: *A string*

Interface name.

slot

Synopsis: *string - one of the following keywords { lm6, lm5, lm4, lm3, lm2, lm1, sm }*

Synopsis: *string*

Line module name of the slot.

Port

Synopsis: *integer*

Synopsis: *string*

Port number on the slot.

Channel Number

Synopsis: *integer*

Synopsis: *string*

Channel number on the port.

state

Synopsis: *string - one of the following keywords { lowerLayerDown, notPresent, dormant, unknown, testing, down, up }*

Status of the interface.

local

Synopsis: *string*

Loacal IP address of the interface.

remote

Synopsis: *string*

Peer IP address.

mask

Synopsis: *string*

Netmask.

20.4.2.3. Frame-relay Statistics

Figure 20.33. Frame Relay Errors Packets Statistics form

The path to this form is *interfaces/wan/t1e1/{line module}/fr-stats/fr-error*.

Frame Length

Synopsis: *unsigned integer*

I-frames not transmitted after a tx. interrupt due to exessive frame length.

Throughput

Synopsis: *unsigned integer*

I-frames not transmitted after a tx. interrupt due to excessive throughput.

Length

Synopsis: *unsigned integer*

Received frames discarded as they were either too short or too long.

Unconfigured DLCI

Synopsis: *unsigned integer*

Discarded I-frames with unconfigured DLCI.

Format Error

Synopsis: *unsigned integer*

Discarded I-frames due to a format error.

No Response

Synopsis: *unsigned integer*

App. didn't respond to the triggered IRQ within the given timeout period.

Signalling Format Error

Synopsis: *unsigned integer*

Discarded In-channel Signalling frames due to a format error.

Invalid Send Sequence

Synopsis: *unsigned integer*

In-channel frames received with invalid Send Seq. Numbers received.

Invalid Receive Sequence

Synopsis: *unsigned integer*

In-channel frames received with invalid Receive Seq. Numbers received.

Unsolicited Response

Synopsis: *unsigned integer*

The number of unsolicited responses from the Access Node.

N391

Synopsis: *unsigned integer*

Timeouts on the T391 timer.

Consecutive T391

Synopsis: *unsigned integer*

Consecutive timeouts on the T391 timer.

N392 Error Threshold

Synopsis: *unsigned integer*

The number of times that the N392 error threshold was reached during N393 monitored events.

Figure 20.34. Frame Relay Controlling Packets Statistics form

The path to the Frame Relay Controlling Packets Statistics form and the Frame Relay Receiving Statistics form is *interfaces/wan/t1e1/{line module}/fr-stats*.

FSE

Synopsis: *unsigned integer*

Full Status Enquiry messages sent.

LIVSE

Synopsis: *unsigned integer*

Link Integrity Verification Status Enquiry messages sent.

FS

Synopsis: *unsigned integer*

Full Status messages received.

LIVS

Synopsis: *unsigned integer*

Link Integrity Verification Status messages received.

CPEI

Synopsis: *unsigned integer*

CPE initializations.

SSEQ

Synopsis: *unsigned integer*

The current Send Sequence Number.

RSEQ

Synopsis: *unsigned integer*

The current Receive Sequence Number.

N392

Synopsis: *unsigned integer*

The current N392 count.

N393

Synopsis: *unsigned integer*

The current N393 count.

Figure 20.35. Frame Relay Receiving Statistics form

Discard

Synopsis: *unsigned integer*

Received I-frames that were discarded due to inactive DLCI.

DEI

Synopsis: *unsigned integer*

I-frames received with the Discard Eligibility (DE) indicator set.

FECN

Synopsis: *unsigned integer*

I-frames received with the FECN bit set.

BECN

Synopsis: *unsigned integer*

I-frames received with the BECN bit set.

20.4.3. Clearing Statistics

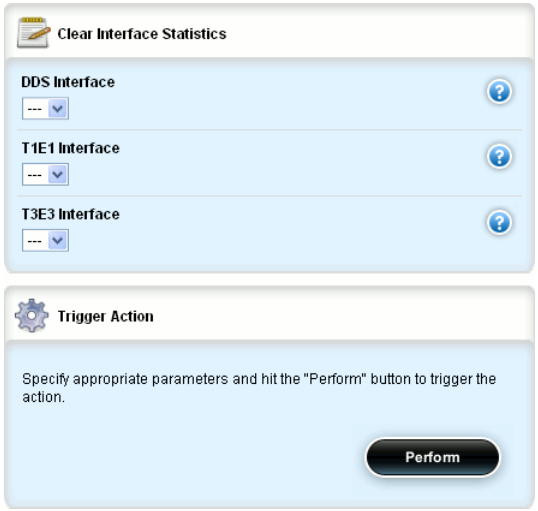

Figure 20.36. Clear Interface Statistics Form And Trigger Action

Statistics can be cleared by specifying the appropriate parameters in the Clear Interface Statistics form and then clicking the *Perform* button on the Trigger Action.

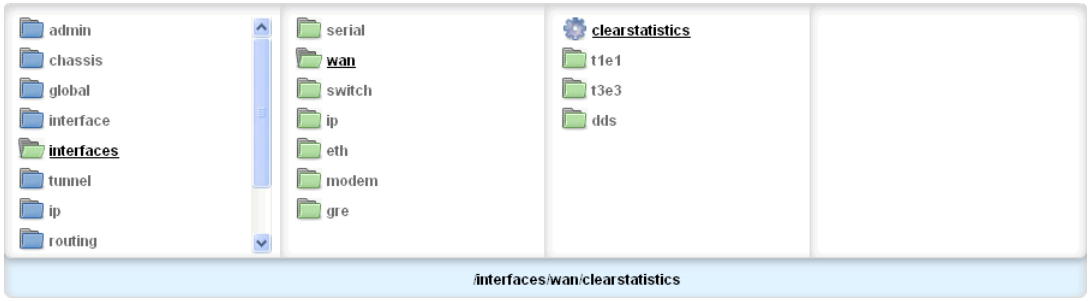

Figure 20.37. Clearstatistics Menu Action

The path to the Clear Statistics forms is *interfaces/wan/clearstatistics*.

20.5. Location Of Interfaces And Labeling

Unlike Ethernet ports, which are found in the same location on all RX1000 units, the location of T1/E1, T3, and DDS ports in your router depends on the number of ports and how they are ordered.

To make labeling easy to understand, all T1/E1, T3/E3, and DDS ports are assigned a unique port number related to the LEDs on the status panel.

20.6. LED Designations

RuggedRouter® includes two sources of LED indicated information about T1/E1 or T3/E3 lines, the T1/E1 or T3/E3 card itself and the LED Panel.

One LED is associated with each line, next to the interface jack. This LED is red when the link is disconnected, flashes green when the link is connecting and remains solid green when the link is established.

The RuggedRouter® also indicates information about T1/E1 or T3/E3 ports on the LED Panel. A pair of LEDs will indicate traffic and link status of the port.

20.7. DDS

Digital Data Services (DDS) is a North American digital transmission method. A DDS line operates synchronously at 56 Kbps over an unloaded 4-wire metallic-pair circuit.

A DDS line is typically a telephone-grade network connection and is often called the “local loop”. A Data Terminal Equipment (DTE) device is attached to the line and transmits data to the telephone company (TELCO), which routes the data to a remote DDS line. A short-haul synchronous-data line driver known as a CSU/DSU terminates the line and attaches to the DTE. The DSU part of the DSU/CSU manages the format of the data signal. The CSU part of the DSU/CSU manages electrical levels and isolation, and provides loopback to the TELCO. A RuggedCom DDS port provides an integrated DTE, DSU, and CSU.

20.7.1. DDS Configuration

To configure DDS, you must first enable the WAN interface supporting DDS. See Section 20.7.1.1, “Enabling the DDS WAN Interface”.

Under */interface/wan{wan slot and port}/dds* you can configure the following:

- set the DDS PPP connection parameters. See Section 20.7.1.2, “Setting DDS PPP Connection Parameters”.
- set the DDS frame relay and DLCI parameters. See Section 20.7.1.3, “Setting DDS Frame Relay Parameters”.

Under interfaces/wan, you can also:

- view and clear DDS statistics. See Section 20.7.2, “Viewing and Clearing DDS Statistics”.
- perform a DDS loopback test. See Section 20.7.1.4, “Performing a DDS Loopback Test”.

20.7.1.1. Enabling the DDS WAN Interface

Before setting DDS parameters, you must first enable the WAN interface supporting DDS.

To enable the WAN DDS interface:

- In edit mode, navigate to */interface/wan{wan slot and port}*.
- On the *Enable Wan Interface* form, select **Enabled**.

Figure 20.38. Enable Wan Interface form

- Commit the changes.

20.7.1.2. Setting DDS PPP Connection Parameters

To set DDS PPP connection parameters, enter edit mode and navigate to */interface/wan{wan slot and port}/dds/connection/ppp*.

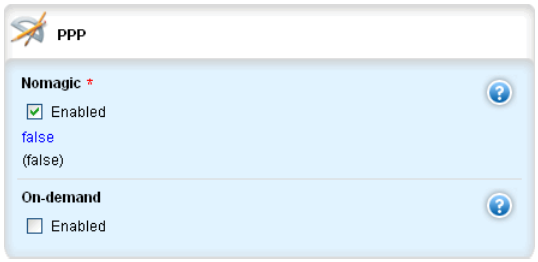

Figure 20.39. PPP form

nomagic

Synopsis: *boolean*

Default: *false*

Disables the Magic Number.

on-demand

This interface is up or down on demand of link fail over.

For more information on the On-demand function, see Section 29.3.4, “Link Backup On Demand”.

20.7.1.3. Setting DDS Frame Relay Parameters

You configure DDS frame relay parameters in two locations:

- */interface/wan{wan slot and port}/dds/connection/framerelay* configures general frame relay parameters
- */interface/wan{wan slot and port}/dds/connection/framerelay/dlci* configures the Data Link Connection Identifier (DLCI) parameters.

To set the DDS frame relay parameters:

- Enter edit mode and navigate to */interface/wan{wan slot and port}/dds/connection/framerelay*.

Figure 20.40. Frame Relay Parameters form

station

Synopsis: *string - one of the following keywords { switch, cpe }*

Default: *cpe*

The behavior of the frame relay connection, i.e. CPE (Customer Premises Equipment) or as a switch.

signal

Synopsis: *string - one of the following keywords { none, q933, lmi, ansi }*

Default: *ansi*

The frame relay link management protocol used.

t391

Synopsis: *integer*

Default: *10*

(Link Integrity Verification polling) Indicates the number of seconds between transmission of in-channel signaling messages. Valid for cpe.

t392

Synopsis: *integer*

Default: *16*

(Verification of polling cycle) Indicates the expected number of seconds between reception of in-channel signaling messages transmitted by cpe. Valid for Switch.

n391

Synopsis: *integer*

Default: *6*

Defines the frequency of transmission of full status enquiry messages. Valid for CPE.

n392

Synopsis: *integer*

Default: *4*

The number of error events (enumerated by n393) for which the channel is declared inactive; valid for either cpe or Switch.

n393

Synopsis: *integer*

Default: *4*

The number of error events on the frame relay channel; valid for either cpe or switch.

eektype

Synopsis: *string - one of the following keywords { request, off }*

Default: *off*

Enables or disables EEK function.

eektimer

Synopsis: *integer*

Default: *5*

The number of seconds to wait before the next EEK message is sent.

To set the DDS frame relay DLCI parameters:

- Enter edit mode and navigate to */interface/wan{wan slot and port}/dds/connection/framerelay*.
- Beside the **framerelay** link, click the ✚ icon and click **<Add dlci>**.
- On the *Key settings* form, enter a number in the range of 16 to 1007 and click **Add**.
- On the *On demand* form, set the On-demand parameter. For more information on the On-demand function, see Section 29.3.4, “Link Backup On Demand”.

20.7.1.4. Performing a DDS Loopback Test

To perform a DDS loopback test, the remote equipment must be able to loop to allow verification of the entire line. If the remote equipment is an RX1000 or RX1500, starting a line loopback will verify both cards and the line. Note that DDS has no standard for performing digital loopback.

To perform a DDS loopback test:

- Navigate to */interfaces/wan/loopback*.
- On the *Loopback Test* form, set the test parameters.

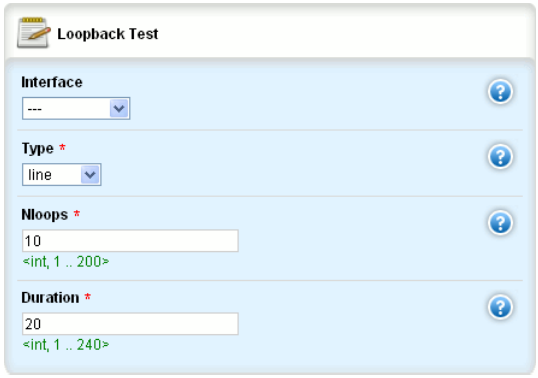

Figure 20.41. Loopback Test form

Interface

Physical interface name.

Type

The loopback type.

Nloops

The number of loops.

Duration

The number of seconds required to run the test.

- On the *Trigger Action* form, click **Perform**.

20.7.2. Viewing and Clearing DDS Statistics

DDS statistics are available when at least one logical interface is configured.

The main DDS statistics menu is available at */interfaces/wan/dds*.

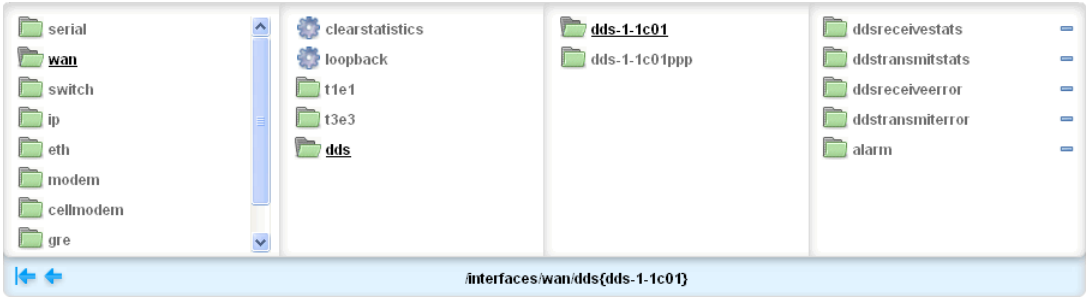

Figure 20.42. DDS Statistics menu

You can view DDS statistics by physical and logical connection. To view DDS physical connection statistics, navigate to:

Path	DDS Physical Connection Statistics
/interfaces/wan/dds{dds physical connection}/ddsreceivestats	Displays DDS physical connection receive statistics.
/interfaces/wan/dds{dds physical connection}/ddstransmitstats	Displays DDS physical connection transmit statistics.
/interfaces/wan/dds{dds physical connection}/ddsreceiveerror	Displays DDS physical connection receive error statistics.
/interfaces/wan/dds{dds physical connection}/ddstransmiterror	Displays DDS physical connection transmit error statistics.
/interfaces/wan/dds{dds physical connection}/alarm	Displays DDS physical connection alarm statistics.

Table 20.1. Paths to DDS Physical Connection Statistics Forms

To view DDS logical connection statistics, navigate to:

Path	DDS Logical Connection Statistics
/interfaces/wan/dds{dds logical connection}/receive	Displays DDS logical connection receive statistics.
/interfaces/wan/dds{dds logical connection}/transmit	Displays DDS logical connection transmit statistics.
/interfaces/wan/dds{dds logical connection}/protocolstats/ppp-stats	Displays DDS PPP protocol statistics.

Table 20.2. Paths to DDS Logical Connection Statistics Forms

20.7.2.1. Clearing DDS Statistics

To clear DDS statistics:

- Navigate to */interfaces/wan/clearstatistics*.
- On the *Clear Interface Statistics* form, set the clear statistics parameters.

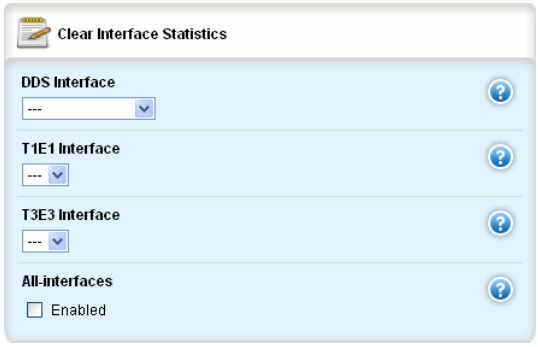

Figure 20.43. Clear Interface Statistics form

DDS Interface

Select the DDS interface for which to clear statistics.

T1E1 Interface

Select the T1E1 interface for which to clear statistics.

T3E3 Interface

Select T3E3 interface for which to clear statistics.

All-interfaces

Clear statistics for all WAN interfaces.

- On the *Trigger Action* form, click **Perform**.

Part III. Routing and Security

This part describes routing and network security:

Tunnelling	Chapter 21, *Tunnelling*
Dynamic Routing	Chapter 22, *Dynamic Routing*
Static Routing	Chapter 23, *Static Routing*
Routing Status	Chapter 24, *Routing Status*
Multicast Routing	Chapter 25, *Multicast Routing*
Firewall	Chapter 26, *Firewall*
Traffic Control	Chapter 27, *Traffic Control*
VRRP	Chapter 28, *VRRP*
LinkFailover	Chapter 29, *Link Failover*

21. Tunnelling

Figure 21.1. Tunnelling menu

The tunnelling menu is accessible from the main menu under *tunnel*. This menu provides access to IPsec, L2TP, L2tunneld and GRE functions.

21.1. IPsec

21.1.1. VPN Fundamentals

IPsec (Internet Protocol SECurity) uses strong cryptography to provide both authentication and encryption services. Authentication ensures that packets are from the right sender and have not been altered in transit. Encryption prevents unauthorized reading of packet contents.

These services allow you to build *secure tunnels through untrusted networks.* Everything passing through the untrusted network is encrypted by the IPsec gateway and decrypted by the gateway at the other end. The result is a Virtual Private Network (VPN), a network which is effectively private even though it includes machines at several different sites connected by the insecure Internet.

The IPsec protocols were developed by the Internet Engineering Task Force (IETF) and are required as part of IP version 6.

Openswan is the open source implementation of IPsec used by ROX™.

The protocols used by IPsec are the Encapsulating Security Payload (ESP) and Internet Key Exchange (IKE) protocols.

ESP provides encryption and authentication (ensuring that a message originated from the expected sender and has not been altered on route).

IKE negotiates connection parameters, including keys, for ESP. IKE is based on the Diffie-Hellman key exchange protocol, which allows two parties without any initial shared secret to create one in a manner immune to eavesdropping.

21.1.1.1. IPsec Modes

IPSec has two basic modes of operation. In transport mode, IPSec headers are added as the original IP datagram is created. The resultant packet is composed of an IP header, IPSec headers and IP payload (including a transport header). Transport mode is most commonly used between IPsec end-stations, or between an end-station and a gateway.

In tunnel mode, the original IP datagram is created normally and then encapsulated into a new IP datagram. The resultant packet is composed of an new IP header, IPSec headers, old IP header and IP payload. Tunnel mode is most commonly used between gateways, the gateway acting as a proxy for the hosts behind it.

21.1.1.2. Policy-Based VPNs

RuggedRouter® supports the creation of policy-based VPNs, which may be characterized as follows:

- IPsec network interfaces are not created.
- The routing table is not involved in directing packets to the IPsec later.
- Only data traffic matching the tunnel's local and remote subnets is forwarded to the tunnel. Normal traffic is routed by one set of firewall rules and VPN traffic is routed based on separate rules.
- The firewall is configured with a VPN zone of type "IPsec".
- As IPsec packets are received, they are decoded, policy-flagged as IPsec-encoded, and presented as having arrived directly via the same network interface on which they were originally received.
- Firewall rules must be written to allow traffic to and from VPN tunnels. These are based on the normal form of source/destination IP addresses and IP protocol and port numbers. These rules, by virtue of the zones they match, use the policy flagging inserted by netkey and route matching data traffic to the proper interface.

21.1.1.3. Supported Encryption Protocols

Openswan supports the following standard encryption protocols:

- 3DES (Triple DES) – Uses three DES encryptions on a single data block, with at least two different keys, to get higher security than is available from a single DES pass. 3DES is the most CPU intensive cipher.
- AES – The Advanced Encryption Standard protocol cipher uses a 128-bit block and 128, 192 or 256-bit keys. This is the most secure protocol in use today, and is much preferred to 3DES due to its efficiency.

21.1.1.4. Public Key And Pre-shared Keys

In *public key cryptography*, keys are created in matched pairs (called public and private keys). The public key is made public while the private key is kept secret. Messages can then be sent by anyone who knows the public key to the holder of the private key. Only the owner of the private key can decrypt the message.

When you want to use this form of encryption, each router configures its VPN connection to use the RSA algorithm and includes the public signature of its peer. The RuggedRouter®'s public signature is available from the output of the *tunnel/ipsec/display-public-key* menu.

In *secret key cryptography*, a single key known to both parties is used for both encryption and decryption.

When you want to use this form of encryption, each router configures its VPN connection to use a secret pre-shared key. The pre-shared key is configured through the *tunnel/ipsec/preshared-key* menu.

21.1.1.5. X509 Certificates

When one side of the VPN connection is placed from a dynamic IP (the so-called "roaming client"), X509 Certificates may be used to authenticate the connection. Certificates are digital signatures that are produced by a trusted source, namely a Certificate Authority (CA). For each host, the CA creates an certificate that contains CA and host information and "signs" the certificate by creating a digest of all the fields in the certificate and encrypting the hash value with its *private key*. The encrypted digest is called a "digital signature". The host's certificate and the CA *public key* are installed on all gateways that the host connects to.

When the gateway receives a connection request, it uses the CA *public key* to decrypt the signature back into the digest. It then recomputes its own digest from the plain text in the certificate and compares

the two. If both digests match, the integrity of the certificate is verified (it was not tampered with), and the public key in the certificate is assumed to be the valid public key of the connecting host.

21.1.1.6. NAT Traversal

Historically, IPSec has presented problems when connections must traverse a firewall providing Network Address Translation (NAT). The Internet Key Exchange (IKE) used in IPSec is not NAT-translatable. When IPSec connections must traverse a firewall IKE messages and IPSec-protected packets must be encapsulated as User Datagram Protocol (UDP) messages. The encapsulation allows the original untranslated packet to be examined by IPSec.

21.1.1.7. Other Configuration Supporting IPSec

If the router is to support a remote IPSec client and the client will be assigned an address in a subnet of a local interface, you must activate proxy ARP for that interface. This will cause the router to respond to ARP requests on behalf of the client and direct traffic to it over its connection.

IPSec relies upon the following protocols and ports:

- protocol 51, IPSEC-AH Authentication Header (RFC2402),
- protocol 50, IPSEC-ESP Encapsulating Security Payload (RFC2046),
- UDP port 500.

You must configure the firewall to accept connections on these ports and protocols. See Section 26.4, “Configuring The Firewall And VPN” in Chapter 26, *Firewall* for details.

21.1.1.8. The Openswan Configuration Process

Each VPN connection has two ends: the local router and the remote router. The Openswan configuration record describing a VPN connection can be used without change at either end. One side of the connection (typically the local side) is designated the “left” side and the other is designated the “right” side.

A convenient method is to configure both ends simultaneously with two command-line interface sessions (or two web browsers) open at the same time. The relevant information is the same in both sessions.

21.1.1.9. IPsec and Router Interfaces

If IPsec works on an interface which could disappear, such as a ppp connection, or if the IP address could change, you need to set the monitor-interface option for the IPsec connection. While this this option is set, IPsec will be restarted when the interface disappears and reappears or the IP address is changed.

For information on setting the monitor-interface option, see the Connection form at *tunnel/ipsec/connection/{line module}*.

21.1.1.10. L2TPD

L2TP stands for “Layer Two Tunneling Protocol”. The main purpose of this protocol is to tunnel PPP packets through an IP network, although it is also able to tunnel other layer 2 protocols.

On RuggedRouter®, L2TPd is used in conjunction with Openswan and PPP to provide support for establishing a secure, private connection with the router using the Microsoft Windows VPN/L2TP client.

L2TPD listens on UDP port 1701. The firewall will need to be configured to allow connections to L2TPD via IPSec but to prevent connections to L2TPD directly without using IPsec.

21.1.2. IPsec Configuration

Figure 21.2. IPsec menu

The IPsec menu is accessible from the main menu under *tunnel*. The path to this menu is *tunnel/ipsec*. The IPsec form appears on the same screen as the *IPsec* menu.

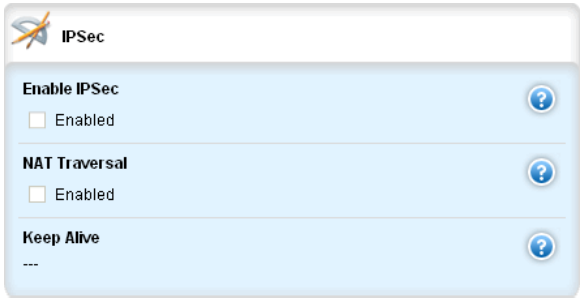

Figure 21.3. IPsec form

The IPsec form is used in configuring IPSec VPN.

Enable IPSec

Enables IPSec.

NAT Traversal

Enables NAT Traversal.

Keep Alive

Synopsis: *unsigned integer*

The delay (in seconds) for sending keepalive packets to prevent NAT router from closing its port when there is not enough traffic on the IPsec connection.

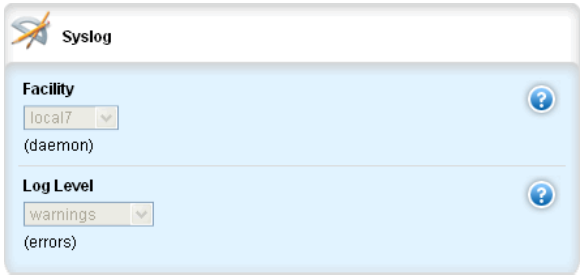

Figure 21.4. Syslog form

The path to the Syslog form is *tunnel/ipsec*.

Facility

Synopsis: *string - one of the following keywords { local7, local6, local5, local4, local3, local2, local1, local0, uucp, user, syslog, news, mark, mail, lpr, kern, daemon, cron, authpriv, auth }*

Default: *daemon*

The log facility.

Log Level

Synopsis: *string - one of the following keywords { warnings, notifications, informational, errors, emergencies, debugging, critical, alerts }*

Default: *errors*

The log level.

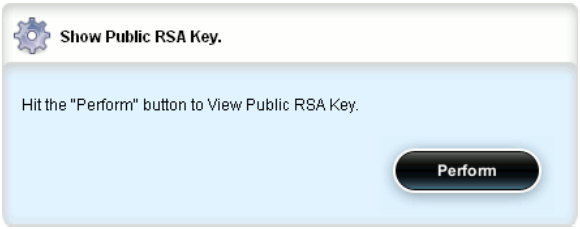

Figure 21.5. Show Public RSA Key form

The path to the Show Public RSA Key form is *tunnel/ipsec/display-public-key*. Clicking on the *Perform* button will display the public key.

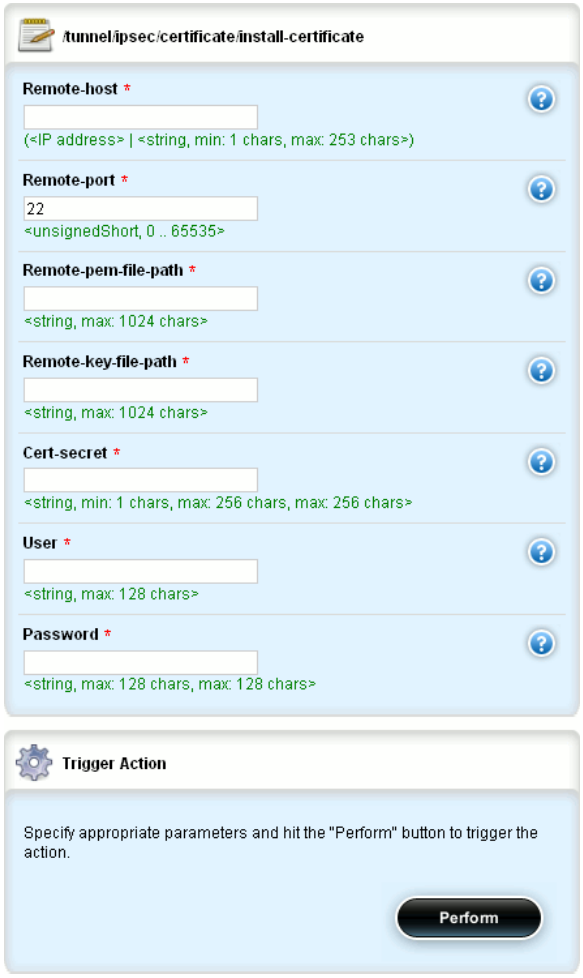

Figure 21.6. Install-Certificate forms

The path to the Install-Certificates forms is *tunnel/ipsec/certificate/install-certificate*. To install the certificates, enter the parameters and then click the *Perform* button.

Figure 21.7. Install-Ca-Certificate forms

The path to the Install-Ca-Certificate forms is *tunnel/ipsec/certificate/install-ca-certificate*. Enter the parameters and then click on the *Perform* button to install the certificates.

Figure 21.8. Install-Crl-File forms

The path to the Install-Crl-File forms is *tunnel/ipsec/certificate/install-crl-file*. To install the files, enter the parameters and then click the *Perform* button.

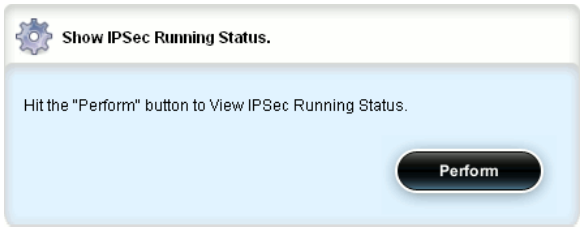

Figure 21.9. Show IPsec Running Status form

The path to the Show IPsec Running Status form is *tunnel/ipsec/status*. To display the IPsec status, click the *Perform* button.

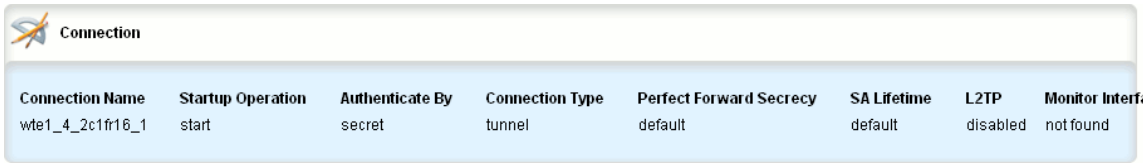

Figure 21.10. Connection table

If data is configured, the path to the Connection table will be *tunnel/ipsec/connection*.

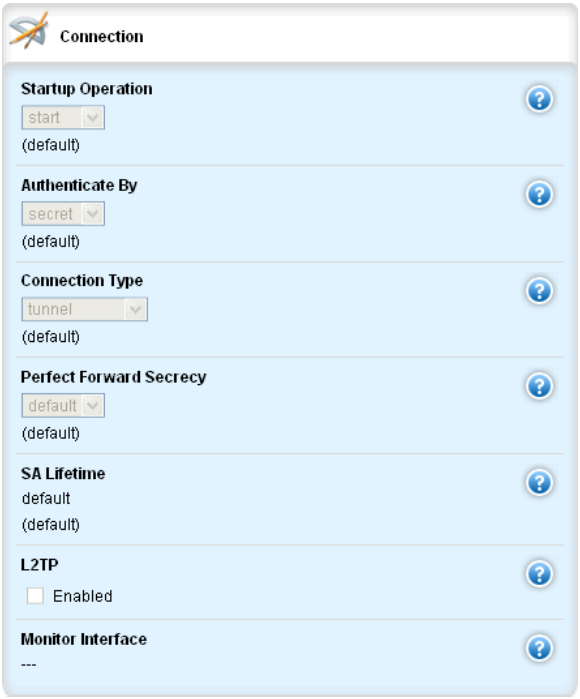

Figure 21.11. Connection form

If data is configured, the path to the Connection form will be *tunnel/ipsec/connection/{line module}*.

The Connection form is used for VPN connection configuration.

Connection Name

Synopsis: *string - the keyword { default }*

Synopsis: *A string conforming to: "[A-Za-z][A-Za-z0-9#%_\-+.,]+"*

The connection name. If the name is the default, this makes it the default setting for all connections.

Startup Operation

Synopsis: *string - one of the following keywords { default, route, start, add, ignore }*

Default: *default*

The action at IPSec startup time.

Authenticate By

Synopsis: *string - one of the following keywords { secret, rsasig, default }*

Default: *default*

The authentication method.

Connection Type

Synopsis: *string - one of the following keywords { default, passthrough, transport, tunnel }*

Default: *default*

The connection type. Tunnel signifies a host-to-host, host-to-subnet, or subnet-to-subnet tunnel; transport signifies a host-to-host transport mode; passthrough signifies that no IPsec processing should be done at all.

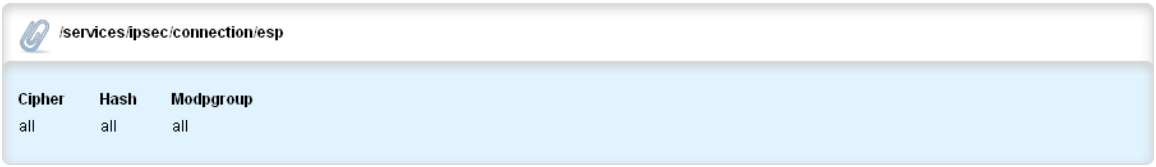

Figure 21.12. ESP table

If data is configured, the path to the ESP table will be *tunnel/ipsec/connection/{line module}/esp*.

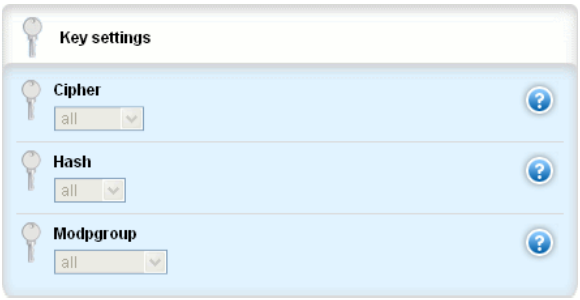

Figure 21.13. ESP Key Settings

If data is configured, the path to the ESP Key Settings form will be to click on *esp/{line module}*.

ESP pertains to the Phase 2 encryption/authentication algorithm to be used for the connection.

Cipher Algorithm

Synopsis: *string - one of the following keywords { any, aes128, aes192, aes256, aes, 3des }*

Cipher algorithm.

Hash Method

Synopsis: *string - one of the following keywords { any, md5, sha1 }*

Hash method.

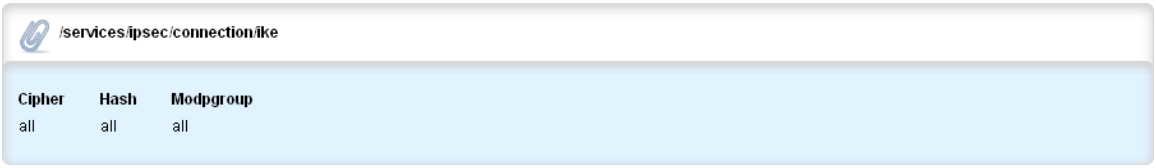

Figure 21.14. IKE table

If data is configured, the path to the IKE table will be *tunnel/ipsec/connection/{line module}/ike*.

IKE pertains to the Phase 1 encryption/authentication algorithm to be used for the connection. A Key Settings Form like the ESP Key Settings Form is also used for IKE.

Cipher Algorithm

Synopsis: *string - one of the following keywords { any, aes128, aes192, aes256, aes, 3des }*

Cipher algorithm.

Hash Method

Synopsis: *string - one of the following keywords { any, md5, sha1 }*

Hash method.

Modpgroup

Synopsis: *string - one of the following keywords { any, modp8192, modp6144, modp4096, modp3072, modp2048, modp1536, modp1024 }*

Modpgroup.

There are right side and left side IPsec forms. The forms for each side are used for IPSec system settings on each side. The forms are the same for both sides, so only the left side forms are shown here.

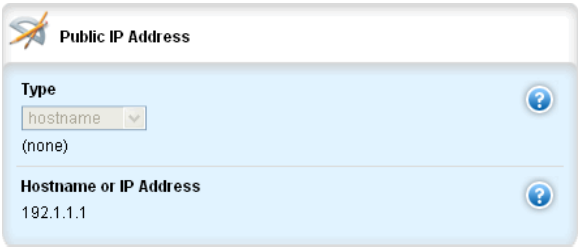

Figure 21.15. Public IP Address form

If data is configured, the path to the Public IP Address form will be *tunnel/ipsec/connection/{line module}/left*. The System Public Key, System Identifier and Nexthop to Other System forms appear on the same screen as the Public IP Address form.

The public-ip is the system identifier.

Type

Synopsis: *string - one of the following keywords { hostname, address, any, default-route, none }*

Default: *none*

Type.

Hostname or IP Address

Synopsis: *A string conforming to: "[^]+"*

Hostname or IP address.

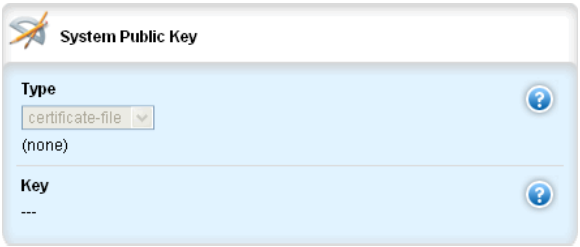

Figure 21.16. System Public Key form

It allows you to select the system's public key.

Type

Synopsis: *string - one of the following keywords { certificate-file, certificate-any, rsasig, none }*

Default: *none*

Type.

Key

Synopsis: *A string conforming to: "[^]+"*

Extra information.

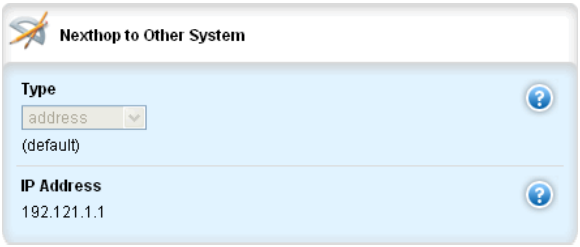

Figure 21.17. Nexthop To Other System form

Type

Synopsis: *string - one of the following keywords { address, default-route, default }*

Default: *default*

Type.

IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

IP address.

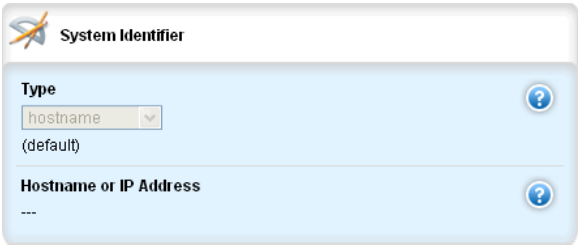

Figure 21.18. System Identifier form

type

Synopsis: *string - one of the following keywords { hostname, address, from-certificate, none, default }*

Default: *default*

Type.

Hostname or IP Address

Synopsis: *A string conforming to: "[^]+"*

Hostname or IP address.

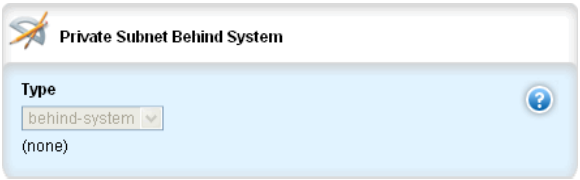

Figure 21.19. Private Subnet Behind System form

If data is configured, the path to the Private Subnet Behind System form will be *tunnel/ipsec/connection/{line module}/left/subnet*.

This form displays the private subnet behind the system.

Type

Synopsis: *string - one of the following keywords { behind-system, none }*

Default: *none*

Type.

Figure 21.20. Network table

The Network table displays a list of subnet addresses.

If data is configured, the path to the Preshared Key table will be *tunnel/ipsec/preshared-key*.

Figure 21.21. Preshared Key table

If data is configured, the path to the Preshared Key form will be *tunnel/ipsec/preshared-key/{line module}*.

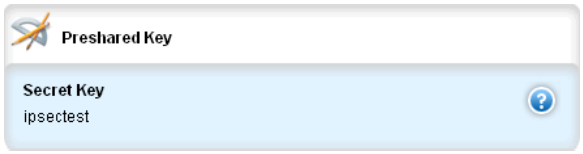

Figure 21.22. Preshared Key form

Remote Address

Synopsis: *string - the keyword { any }*

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *A string conforming to: "[^]+"*

The remote address.

Local Address

Synopsis: *string - the keyword { any }*

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *A string conforming to: "[^]+"*

The local address.

Secret Key

Synopsis: *"AES CFB128"-encrypted string*

The pre-shared key.

21.2. L2TP Tunnelling Configuration

Figure 21.23. L2TP menu

The path to the L2TP menu is *tunnel/l2tp*. The L2TP, DNS Server, PPP Options and WINS server forms appear on the same screen as this menu.

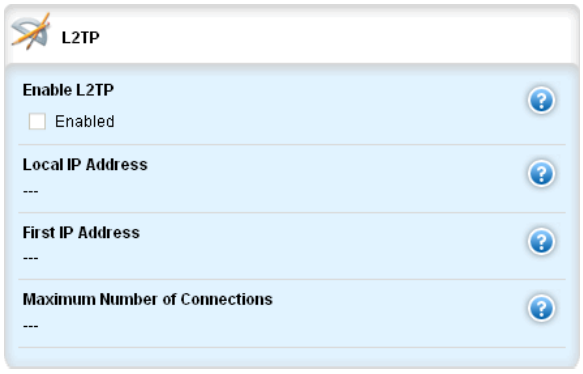

Figure 21.24. L2TP form

Enable L2TP

Enable L2TP.

Local IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

Local IP address.

First IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

First address in remote IP address pool.

Maximum Number of Connections

Synopsis: *unsigned integer*

Maximum number of connections.

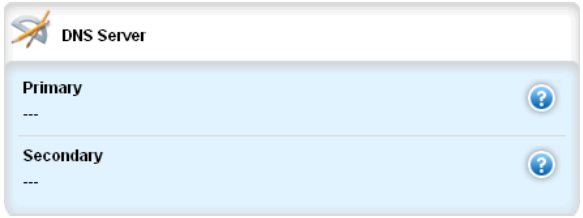

Figure 21.25. DNS Server form

Primary

Synopsis: *IPv4 address in dotted-decimal notation*

Primary DNS server.

Secondary

Synopsis: *IPv4 address in dotted-decimal notation*

Secondary DNS server.

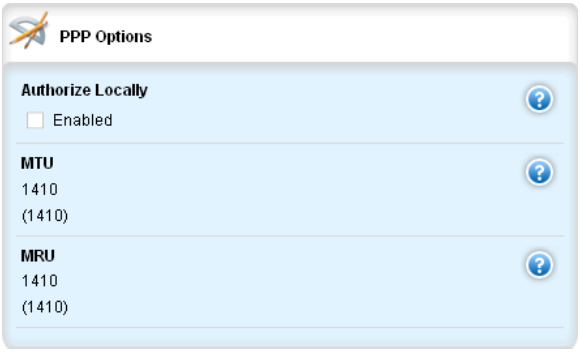

Figure 21.26. PPP Options form

Before enabling the Authorize Locally field on the PPP Options form, you need to add a PPP user name and password under the global/ppp/profiles/dialin menu. If you are not enabling the Authorize Locally field, you need to configure the Radius server for ppp authentication under the global/ppp/radius menu. For more information on PPP, see Chapter 13, PPP Users or Section 18.2.1.5, "V90".

Authorize Locally

Authorize locally instead of using radius server.

MTU

Synopsis: *integer*

Default: *1410*

Maximum Transmit Unit (MTU) or maximum packet size transmitted.

MRU

Synopsis: *integer*

Default: *1410*

Maximum Receive Unit (MRU) or maximum packet size passed when received.

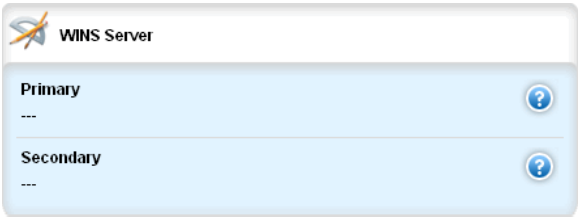

Figure 21.27. WINS Server form

Primary

Synopsis: *IPv4 address in dotted-decimal notation*

Primary WINS server.

Secondary

Synopsis: *IPv4 address in dotted-decimal notation*

Secondary WINS server.

21.3. Layer 2 Tunnelling

RuggedRouter® is capable of extending the range of services that communicate solely via Layer 2 protocols (i.e. at the level of Ethernet) by tunneling them over routed IP networks. The Layer 2 Tunnel Daemon supports the IEC61850 GOOSE protocol as well as a generic mechanism for tunneling by Ethernet type.

You can configure GOOSE tunnels and generic Layer 2 tunnels and also view tunnel status and statistics.

21.3.1. IEC61850 GOOSE Fundamentals

IEC61850 is an international standard for substation automation. It is a part of the International Electrotechnical Commission's (IEC) Technical Committee 57 (TC57) architecture for electric power systems. An important feature of IEC61850 is the fast transfer of event data. Transfers of Generic Substation Events (GSEs) are accomplished through the GOOSE (Generic Object Oriented Substation Event) protocol.

IEC61850 uses Layer 2 multicast frames to distribute its messages and hence is incapable of operating outside of a switched Ethernet Network. The GOOSE tunnel feature provides a capability to bridge GOOSE frames over a WAN.

GOOSE tunnels provide the following features:

- GOOSE traffic is bridged over the WAN via UDP/IP.
- One GOOSE traffic source can be mapped to multiple remote router Ethernet interfaces in mesh fashion.
- To reduce bandwidth consumption, GOOSE daemons may be located at each of the "legs" and at the center of a star network. The centrally located daemon will accept GOOSE packets and re-distribute them.
- Statistics report availability of remote GOOSE daemons, packet counts and Round Trip Time (RTT) for each remote daemon.
- When Virtual Router Redundancy Protocol (VRRP) is employed, GOOSE transport is improved by sending redundant GOOSE packets from each VRRP gateway.
- You can enable GOOSE forwarding by configuring a generic Layer 2 tunnel. When configured, RuggedRouter® listens for GOOSE packets on one VLAN and forwards them to another VLAN.

21.3.1.1. GOOSE Tunnel Implementation Details

The GOOSE protocol is supported by the Layer 2 Tunnel Daemon. The daemon listens to configured Ethernet interfaces and to the network itself (i.e. for tunnel connections from other daemon instances) on a configurable UDP port.

The Media Access Control (MAC) destination address of frames received from Ethernet is inspected in order to determine which GOOSE group they are in. The frames are then encapsulated in network headers and forwarded (with MAC source and destination addresses intact) to the network as GOOSE packets.

IEC61850 recommends that the MAC destination address should be in the range 01:0c:cd:01:00:00 to 01:0c:cd:01:01:ff.

GOOSE Packets received from the network are stripped of their network headers and forwarded to Ethernet ports configured for the same multicast address. The forwarded frames contain the MAC source address or the originating device, and not that of the transmitting interface. The VLAN used will be that programmed locally for the interface and may differ from the original VLAN. The frame will be transmitted with the highest 802.1p priority level (p4).

Packets received from the network will also be forwarded to any other remote daemons included in the group.

To enable forwarding for GOOSE packets, configure a generic Layer 2 tunnel to listen for GOOSE packets on one VLAN and forward them to a second VLAN. To configure the generic Layer 2 tunnel for this operation, set the following for the tunnel:

- **Ethernet Interface**: select the VLAN on which the GOOSE packets orginate.
- **Ethernet Type**: set as 0x8868. 0x8868
- **Remote Daemon**: select the VLAN to which to forward the GOOSE packets.

21.3.2. Generic Layer 2 Tunnel Fundamentals

The Layer 2 Tunnel Daemon also supports a generic mode of operation based on the Ethernet type of Layer 2 data traffic seen by the router. Multiple tunnels may be configured, each one with:

- Ethernet type
- Tunnel ingress (Ethernet interface)
- Tunnel egress (either another locally connected Ethernet interface, or the remote IP address of another Layer 2 Tunnel daemon instance running on another RuggedRouter®)

21.3.2.1. Generic Tunnel Implementation Details

For each tunnel configured, the daemon monitors the specified Ethernet interface for Ethernet (Layer 2) frames of the specified type. If the configured egress is another local Ethernet port, frames are simply forwarded on that port, unmodified.

If the configured tunnel egress is a remote IP address, the daemon encapsulates the frames and forwards them to that address, where a corresponding Layer 2 Tunnel Daemon must be configured to receive tunneled frames for local retransmission. Encapsulation headers are stripped in order that the retransmitted frames are identical to those received at the tunnel ingress.

Other notes:

- Source and destination Ethernet MAC addresses are preserved, whether they are forwarded locally or remotely.
- Packets received from the network will also be forwarded to any other remote daemons included in the group.
- The UDP port number for inter-daemon communication must be the same throughout the network
- Enabling Generic L2 Tunneling on an Ethernet interface does not interfere with other (Layer 3) networking configuration on that interface, e.g. firewall rules, IP routing, etc.

Avoid network configurations where the daemons can form a traffic loop. The simplest such configuration is a triangle network where each daemon forwards to two other routers. Frames arriving at one router will start cycling in clockwise and counterclockwise directions.

To avoid such “packet storms”, frames forwarded to the network are tagged with an initial time to live count. The count is decremented at each relay to the network and prevents the frame from being relayed indefinitely.

21.3.3. Layer 2 Tunnelling Configuration

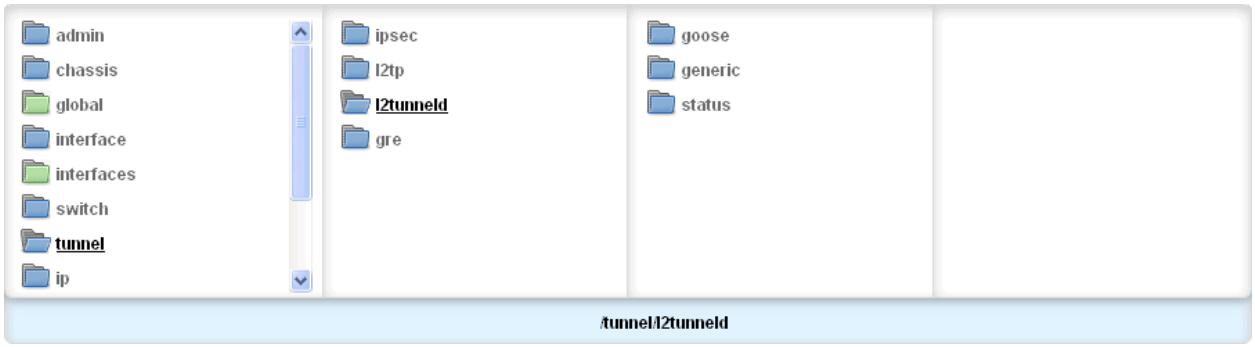

Figure 21.28. L2tunneld menu

The path to the L2tunneld (Layer 2) menu is *tunnel/l2tunneld*. The L2 Tunnel Daemon form appears on the same screen as this menu.

Figure 21.29. L2 Tunnel Daemon form

This form configures general settings for the daemon that apply to all supported tunnel configurations. The UDP-port field configures the port used by the daemon to communicate with other daemons. The Beacon-interval field configures how often a Round Trip Time (RTT) measurement message is sent to each remote peer. The interval takes the values “Off” to disable RTT measurement or a time of 10 – 3600 seconds.

All Layer 2 Tunnel daemons in the network must use the same port number. If the router employs a firewall, ensure that a hole is opened for each of the remote daemons using this port number.

enabled

Enables the L2 protocols server.

udp-port

Synopsis: *integer*

Default: *1311*

The UDP port to communicate with other deamon

beacon-interval

Synopsis: *string - the keyword { off }*

Synopsis: *integer*

Default: *60*

The Round Trip Time (RTT) of sending message

21.3.3.1. Goose

The forms and tables in this section are located under *tunnel/l2tunneld/goose*.

Figure 21.30. Goose Tunnel table

This table displays configured GOOSE tunnels.

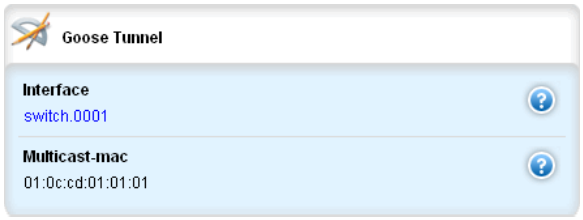

Figure 21.31. Goose Tunnel form

name

Synopsis: *A string*

Description of goose tunnel

interface

Synopsis: *A string*

The interface to listen on for goose frames

multicast-mac

Synopsis: *Multicast Ethernet MAC address in colon-separated hexadecimal notation*

The multicast MAC address to listen for

Figure 21.32. Remote Daemon of Goose Tunnel table

ip-address

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address of remote L2 protocol server

21.3.3.2. Generic

The forms and tables in this section are located under *tunnel/l2tunneld/generic*.

Figure 21.33. Generic L2 Tunnel table

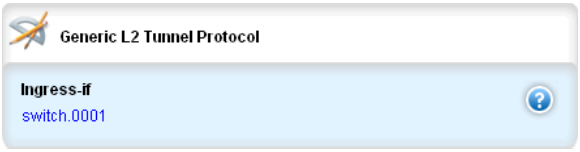

Figure 21.34. Generic L2 Tunnel Protocol form

name

Synopsis: *A string*

Description of goose tunnel

ingress-if

Synopsis: *A string*

The interface to listen on for Ethernet type frames

Figure 21.35. Generic L2 Tunnel Egress Interface table

egress-if

Synopsis: *A string*

Egress interface for Ethernet type frames

Figure 21.36. L2 Ethernet Type table

type

Synopsis: *string - the keyword { iso }*

Synopsis: *A string conforming to: "(0x[0-9A-Fa-f]{4})"*

Ethernet type to be forwarded (ie. 0xFEFE)

21.3.3.3. Status

The forms and tables in this section are located under *tunnel/l2tunneld/status/*. The submenus Goose, Generic and Round-Trip-Time are accessible from the *status* menu.

21.3.3.3.1. Status Goose

The forms and tables in this section are located under *tunnel/l2tunneld/status/goose*.

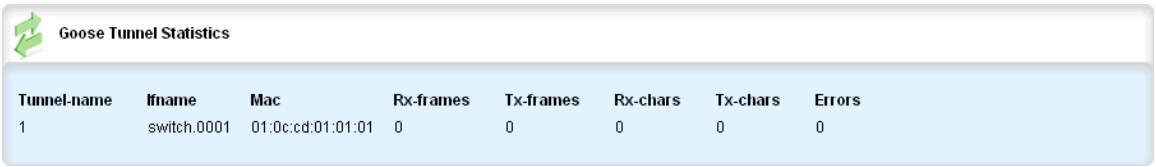

Figure 21.37. Goose Tunnel Statistics table

Figure 21.38. Goose Tunnel Statistics form

tunnel-name

Synopsis: *A string*

Goose Tunnel name

ifname

Synopsis: *A string*

VLAN Interface name

mac

Synopsis: *Multicast Ethernet MAC address in colon-separated hexadecimal notation*

Multicast Destination MAC Address of Goose message

rx-frames

Synopsis: *unsigned integer*

The number of frames received over the tunnel

tx-frames

Synopsis: *unsigned integer*

The number of frames transmitted over the tunnel

rx-chars

Synopsis: *unsigned integer*

The number of bytes received over the tunnel

tx-chars

Synopsis: *unsigned integer*

The number of bytes transmitted over the tunnel

errors

Synopsis: *unsigned integer*

The number of errors over the tunnel

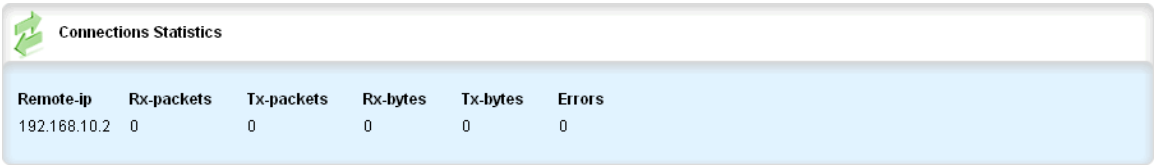

Figure 21.39. Connections Statistics table

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of remote goose daemon

rx-packets

Synopsis: *unsigned integer*

The number of frames received over the tunnel

tx-packets

Synopsis: *unsigned integer*

The number of frames transmitted over the tunnel

rx-bytes

Synopsis: *unsigned integer*

The number of bytes received over the tunnel

tx-bytes

Synopsis: *unsigned integer*

The number of bytes transmitted over the tunnel

errors

Synopsis: *unsigned integer*

The number of errors over the tunnel

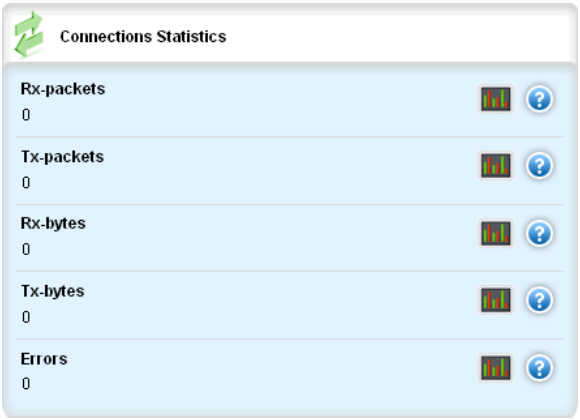

Figure 21.40. Connections Statistics form

21.3.3.3.2. Status Generic

The forms and tables in this section are located under *tunnel/l2tunneld/status/generic*.

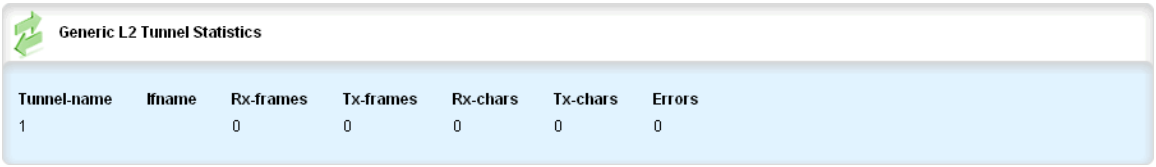

Figure 21.41. Generic L2 Tunnel Statistics table

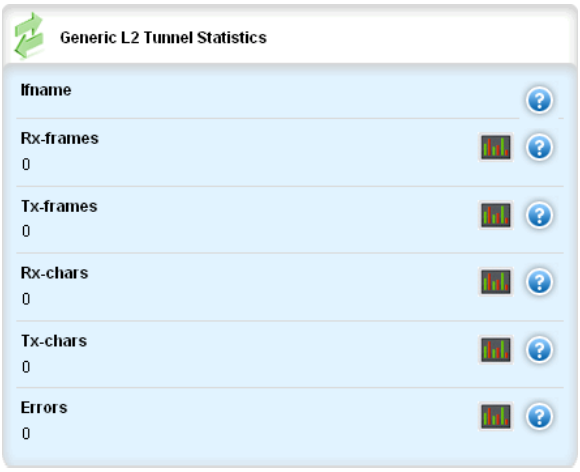

Figure 21.42. Generic L2 Tunnel Statistics form

tunnel-name

Synopsis: *A string*

Goose Tunnel name

ifname

Synopsis: *A string*

VLAN Interface name

rx-frames

Synopsis: *unsigned integer*

The number of frames received over the tunnel

tx-frames

Synopsis: *unsigned integer*

The number of frames transmitted over the tunnel

rx-chars

Synopsis: *unsigned integer*

The number of bytes received over the tunnel

tx-chars

Synopsis: *unsigned integer*

The number of bytes transmitted over the tunnel

errors

Synopsis: *unsigned integer*

The number of errors over the tunnel

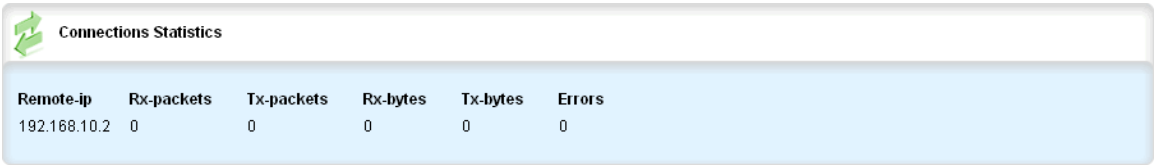

Figure 21.43. Connections Statistics table

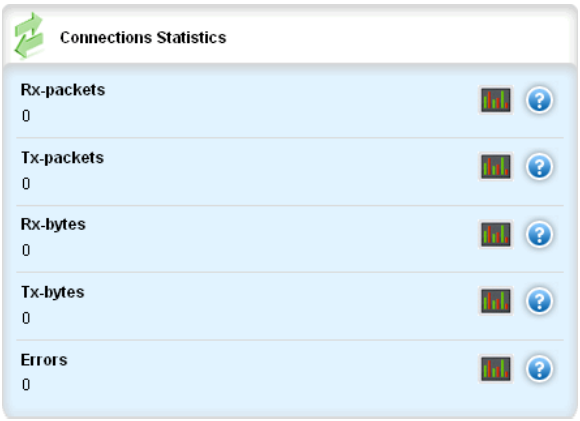

Figure 21.44. Connections Statistics form

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of remote goose daemon

rx-packets

Synopsis: *unsigned integer*

The number of frames received over the tunnel

tx-packets

Synopsis: *unsigned integer*

The number of frames transmitted over the tunnel

rx-bytes

Synopsis: *unsigned integer*

The number of bytes received over the tunnel

tx-bytes

Synopsis: *unsigned integer*

The number of bytes transmitted over the tunnel

errors

Synopsis: *unsigned integer*

The number of errors over the tunnel

21.3.3.3.3. Status Round-Trip-Time

The forms and tables in this section are located under *tunnel/l2tunneld/status/round-trip-time*.

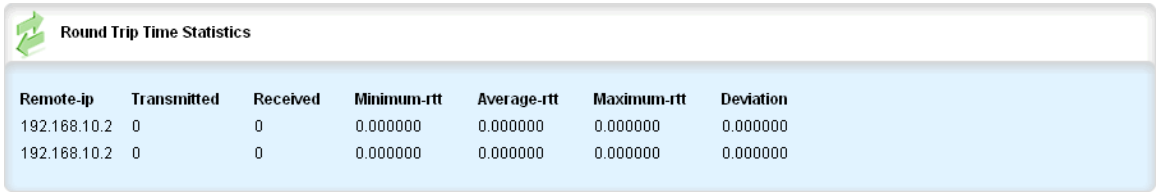
Round Trip Time Statistics
Remote-ip Transmitted Received Minimum-rtt Average-rtt Maximum-rtt Deviation
192.168.10.2 0 0 0.000000 0.000000 0.000000 0.000000
192.168.10.2 0 0 0.000000 0.000000 0.000000 0.000000

Figure 21.45. Round Trip Time Statistics table

The Round Trip Time Statistics table reflects the measured RTT to each remote daemon. The minimum, average, maximum and standard deviation of times is presented. Entries with a large difference between the Transmitted and Received fields indicate potential problems.

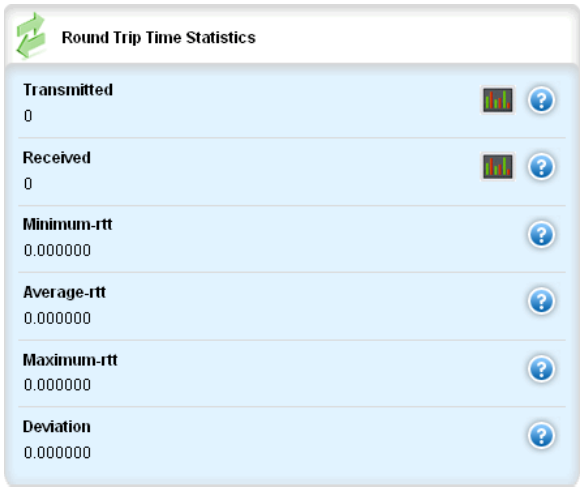

Figure 21.46. Round Trip Time Statistics form

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of remote goose daemon

transmitted

Synopsis: *unsigned integer*

The number of beacon frames transmitted over the tunnel

received

Synopsis: *unsigned integer*

The number of beacon frames transmitted over the tunnel

minimum-rtt

Synopsis: *A string*

The Minimum Beacon Round-Trip-Time

average-rtt

Synopsis: *A string*

The Average Beacon Round-Trip-Time

maximum-rtt

Synopsis: *A string*

The Maximum Beacon Round-Trip-Time

deviation

Synopsis: *A string*

The Standard Deviation

21.4. Generic Routing Encapsulation (GRE)

ROX™ is able to encapsulate multicast traffic and IPv6 packets and transport them through an IPv4 network tunnel.

A GRE tunnel can transport traffic through any number of intermediate networks. The key parameters for GRE in each router are the tunnel name, local router address, remote router address and remote subnet.

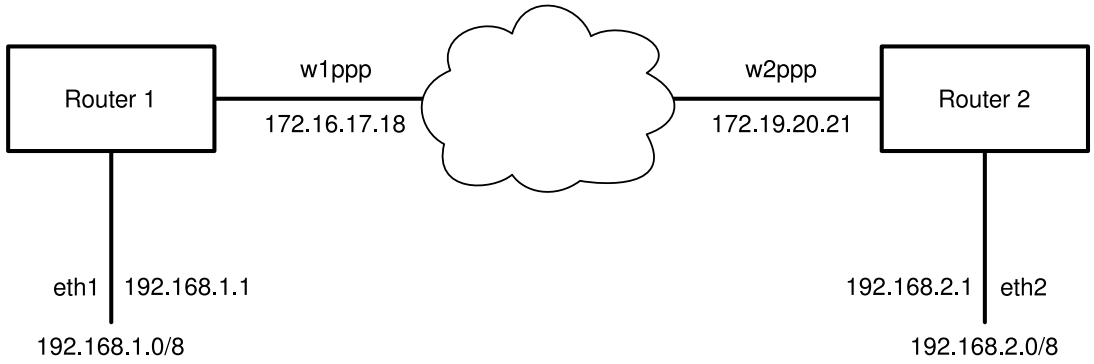

Figure 21.47. GRE Example

In the above example, Router 1 will use a GRE tunnel with a local router address of 172.16.17.18, a remote router address of 172.19.20.21, and a remote subnet of 192.168.2.0/24.

If you are connecting to a CISCO router (in place of Router 1 in the example above), the local router address corresponds to the CISCO IOS "source" address and the remote router address corresponds to the "destination" address.

You may also set a cost for the tunnel. If another method of routing between Router1 and Router2 becomes available, the tunnelled packets will flow through the lowest cost route. Optionally, you can restrict the packets by specifying the local egress device (in the case of router1, w1ppp).

21.4.1. Generic Routing Encapsulation Configuration

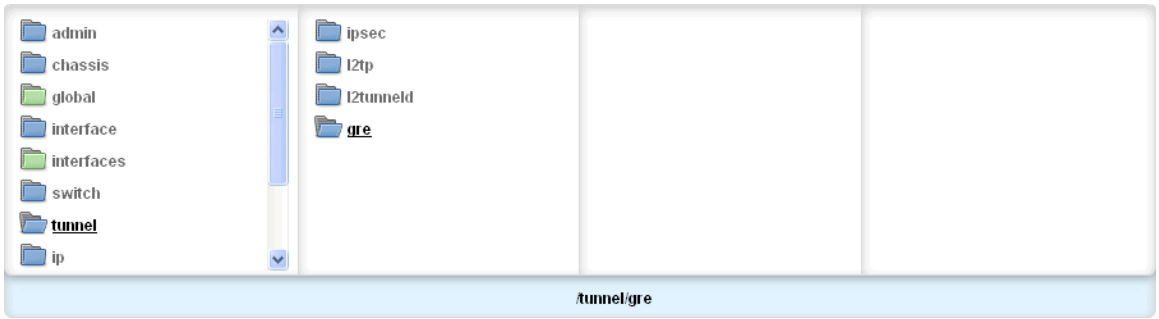

Figure 21.48. Generic Routing Encapsulation (GRE) menu

The path to this menu is *tunnel/gre*. If GRE tunnels are configured, they will be accessible via linked submenus.

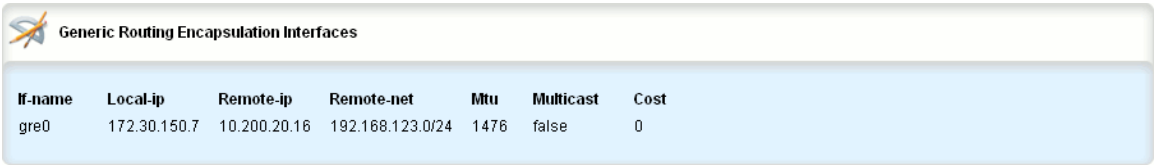
Generic Routing Encapsulation Interfaces
If-name Local-ip Remote-ip Remote-net Mtu Multicast Cost
gre0 172.30.150.7 10.200.20.16 192.168.123.0/24 1476 false 0

Figure 21.49. Generic Routing Encapsulation Interfaces table

The Generic Routing Encapsulation Interfaces table appears on the same screen as the *GRE* menu.

Figure 21.50. Generic Routing Encapsulation Interfaces form

The path to the Generic Routing Encapsulation Interfaces form is *tunnel/gre* and then clicking on one of the linked submenus that follow (for example, gre0).

if-name

Synopsis: *A string conforming to: "[A-Za-z]{1}[-0-9A-Za-z]{0,9}.*"*

The GRE tunnel network interface name. The prefix 'gre-' will be added to this interface name.

local-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address of the local end of the tunnel

remote-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The IP address of the remote end of the tunnel

remote-net

Synopsis: *IPv4 address and prefix in CIDR notation*

The target network of the remote end of the tunnel (xxx.xxx.xxx.xxx/xx)

mtu

Synopsis: *integer*

Default: *1476*

The MTU of the GRE interface

multicast

Enables multicast traffic on the tunnel interface

cost

Synopsis: *integer*

Default:

The routing cost associated with networking routing that directs traffic through the tunnel

22. Dynamic Routing

22.1. Introduction

This chapter familiarizes the user with:

- Enabling the Dynamic Routing Suite
- Enabling and starting OSPF, RIP, and BGP
- Configuring OSPF, RIP, and BGP
- Obtaining OSPF, RIP, and BGP Status
- OSPF and VRRP

22.1.1. RIP, OSPF, and BGP

Dynamic routing is provided by a set of routing protocol daemons. The following daemons manage routing: core, ripd, ospfd, and bgpd.

The core daemon handles interfacing with the kernel to maintain the router's routing table and to check link statuses. It tells RIP, OSPF, and BGP what state links are in, what routes are in the routing table, and some information about the interfaces.

The ripd, ospfd, and bgpd daemons handle communications with other routers using the RIPv2, OSPFv2, and BGP protocols respectively, and decide which routers are preferred to forward to for each network route known to the router.

In complex legacy networks, RIP, OSPF, and BGP may be active on the same router at the same time. Typically, however, one of them is employed.

22.1.2. RIP Fundamentals

The Routing Information Protocol determines the best path for routing IP traffic over a TCP/IP network based on the number of hops between any two routers. It uses the shortest route available to a given network as the route to use for sending packets to that network.

The ROX™ RIP daemon (ripd) is an RFC1058 compliant implementation of RIP support RIP version 1 and 2. RIP version 1 is limited to obsolete class based networks, while RIP version 2 supports subnet masks as well as simple authentication for controlling which routers to accept route exchanges with.

RIP uses network and neighbor entries to control which routers it will exchange routes with. A network is either a subnet or a physical (broadcast-capable) network interface. Any router that is part of that subnet or connected to that interface may exchange routes. A neighbor is a specific router to exchange routes with specified by its IP address. For point to point links (T1/E1 links for example) one must use neighbor entries to add other routers to exchange routes with. The maximum number of hops between two points on a RIP network is 15, placing a limit on network size.

Link failures will eventually be noticed although it is not unusual for RIP to take many minutes for a dead route to disappear from the whole network. Large RIP networks could take over an hour to converge when link or route changes occur. For fast convergence and recovery, OSPF is a much better choice. RIP is a fairly old routing protocol and has mostly been superseded by OSPF.

22.1.3. OSPF Fundamentals

The Open Path Shortest First (OSPF) protocol routing determines the best path for routing IP traffic over a TCP/IP network based on link cost and quality. Unlike static routing, OSPF takes link failures and other network topology changes into account. Unlike the RIP routing protocol, OSPF provides less router to router update traffic.

The ROX™ OSPF daemon (ospfd) is an RFC 2178 compliant implementation of OSPFv2. The daemon also adheres to the RFC2370 (Opaque LSA) and RFC3509 (ABR-Types) extensions.

OSPF network design usually involves partitioning a network into a number of self-contained areas. The areas are chosen to minimize intra-area router traffic, making more manageable and reducing the number of advertised routes. Area numbers are assigned to each area. All routers in the area are known as Area routers. If traffic must flow between two areas a router with links in each area is selected to be an Area Border router, and serves as a gateway.

22.1.3.1. Link State Advertisements

When an OSPF configured router starts operating it issues a hello packet. Routers having the same OSPF Area, hello-interval and dead-interval timers will communicate with each other and are said to be neighbors

After discovering its neighbors, a router will exchange Link State Advertisements in order to determine the network topology.

Every 30 minutes (by default), the entire topology of the network must be sent to all routers in an area. If the link speeds are too low, the links too busy or there are too many routes, then some routes may fail to get re-announced and will be aged out.

Splitting the network into smaller areas to reduce the number of routes within an area or reducing the number of routes to be advertised may help to avoid this problem.

In shared access networks (i.e. routers connected by switches or hubs) a designated router and a backup designated are elected to receive route changes from subnets in the area. Once a designated router is picked, all routing state changes are sent to the designated router, which then sends the resulting changes to all the routers.

The election is decided based on the priority assigned to the interface of each router. The highest priority wins. If the priority is tied, the highest router-id wins.

22.1.4. Key OSPF And RIP Parameters

22.1.4.1. Network Areas

Network areas determine the regions within which routes are distributed to other routers. The subnets at a particular router can be added to its OSPF Area. The router will advertise these subnets to all routers in its area.

OSPF areas must be designed such that no single link failure will cause the network to be split into two disjoint networks.

A router can be part of multiple areas and function as a gateway between areas. When multiple areas are used on a network, area 0 is the backbone area. All areas must have a router connecting them to area 0.

22.1.4.2. Router-ID

Defines the ID of the router. By default this is the highest IP assigned to the router. It is often a good idea to configure this value manually to avoid the router-id changing if interfaces are added or deleted from the router. During elections for designated router, the router-id is one of the values used to pick the winner. Keeping the router-id fixed will avoid any unexpected changes in the election of the master router.

22.1.4.3. Hello Interval and Dead Interval

The hello interval is the time between transmission of OSPF Hello packets. The dead interval is the time to wait without seeing an OSPF Hello packet before declaring a neighboring router dead and discarding its routes. It is recommended that the dead interval be at least four times the hello interval for reliable operation.

Lower values of these settings will help to speed up the change in network routes when the topology of the network changes. It will also increase the load on the router and the links, due to higher traffic caused by the increase in messages. Lower values will also put limits on the number of routes that can be distributed within an area, as will running over slower links.

OSPF will not work properly if the Hello Interval and Dead Interval are not identical on every router in an area.

22.1.4.4. Active/Passive Interface Default

OSPF regards router interfaces as either passive or active, sending OSPF messages on active interfaces and ignoring passive interfaces. By default, newly created interfaces are viewed as passive from OSPF until they are configured active. This is more efficient and secure for the router. The default type for new interfaces is controlled by the passive interface default option in the OSPF Global Parameters.

The default setting of Passive Interface Default means that you must explicitly configure interfaces active before OSPF will attempt to use them.

22.1.4.5. Redistributing Routes

Routes for subnets which are directly connected to the router but are not part of the OSPF area or RIP or BGP networks can be advertised if "redistribute connected" is enabled in the OSPF, RIP, or BGP Global Parameters. Static routes and routes handled by the kernel can also be redistributed if "redistribute static" and "redistribute kernel" are enabled, respectively.

22.1.4.6. Link Detect

Link detection is enabled for active network interfaces. Link detection ensures that the appropriate routing daemon is notified when an interface goes down and stops advertising subnets associated with that interface. The routing daemon resumes advertising the subnet when the link is restored. This allows routing daemons to detect link failures more rapidly, as the router does not have to wait for a “dead interval” to time out). Link detection also causes ”redistributed“ routes to start and stop being advertised based on the status of their interface links.

22.1.4.7. Configuring OSPF Link Costs

Link cost determines which route to use when multiple links can reach a given destination. By default, OSPF assigns the same cost to all links unless it is provided with extra information about the links. Each interface is assumed to be 10Mbit unless specified otherwise in the auto-cost bandwidth setting found under the *ip* menu.

The default OSPF reference bandwidth for link cost calculations is 100Mbit. The reference bandwidth divided by the link bandwidth gives the default cost for a link,which by default is 10. If a specific bandwidth is assigned to each link, the costs take this into account.

You can manually assign a link cost in the OSPF Interface Configuration for each interface. Do this for cases where the speed of the link should not be used as the method for choosing the best link.

22.1.4.8. OSPF Authentication

OSPF authentication is used when it is desirable to prevent unauthorized routers from joining the OSPF network. By enabling authentication and configuring a shared key on all the routers, only routers which have the same authentication key will be able to send and receive advertisements within the OSPF network. Authentication adds a small overhead due to the encryption of messages, so is not to be preferred on completely private networks with controlled access.

22.1.4.9. RIP Authentication

RIP authentication is used when it is desirable to prevent unauthorized routers from joining the network. RIP authentication is supported by per-interface configuration or the use of key-chains. Separate key chains spanning different groups of interfaces and having separate lifespans are possible. By enabling authentication and configuring a shared key on all the routers, only routers which have the same authentication key will be able to send and receive advertisements within the RIP network.

22.1.4.10. Administrative Distances

The router may work with different routing protocols at the same time, as well as employing local interface and statically assigned routes. An administrative distance, (from 0 to 255) is a rating of the trustworthiness of a routing information source. For a given route, the protocol having the lowest administrative distance will be chosen. By default the distances for a connected interface is 0 and for a static route is 1. By default, OSPF will set an administrative distance of 110 and RIP will set a distance of 120.

22.1.5. OSPF And VRRP Example Network

This network consists of three routers connected in a ring with T1/E1 links. Router 1 and 2 and the switched network represent a remote site in which the routers supply a redundant gateway to the hosts via VRRP and the T1/E1 links supply a redundant network connection to the rest of the network.

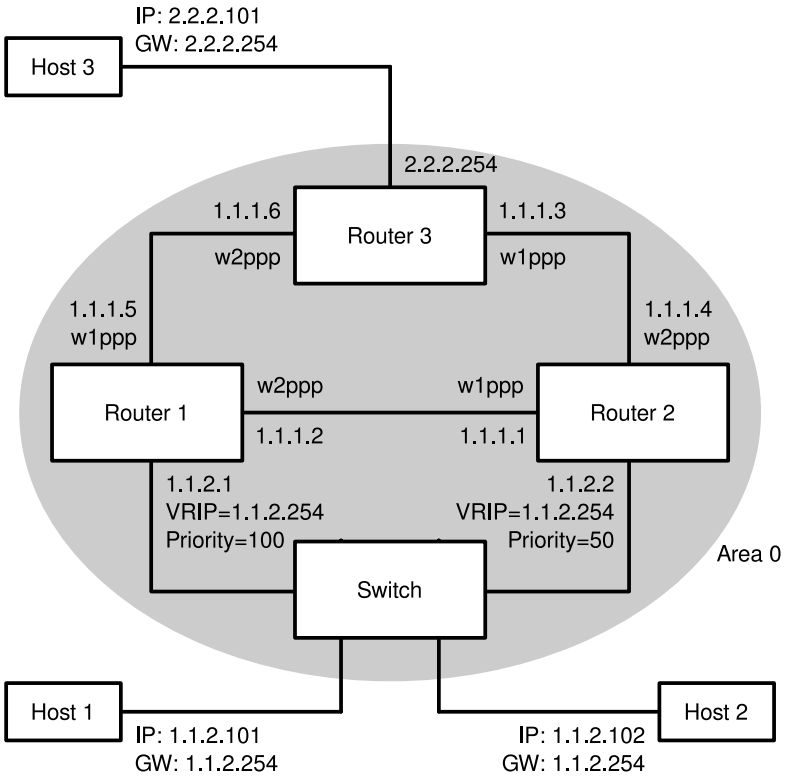

Figure 22.1. OSPF and VRRP Example

22.1.5.1. Area And Subnets

As the OSPF design is simple, an area of 0 is used. The three point-to-point T1/E1 links are placed in the area by adding 1.1.1.0/24 to it. Router 1 and 2 will include their Ethernet links by adding subnet 1.1.2.0/24 to their area descriptions. Router 3 must also include 2.2.2.0/24 in its area description so that its existence is advertised.

The point-to-point T1/E1 interfaces and Ethernet interfaces on Router 1 and 2 must be made active. The Ethernet interface on Router 3 can be left passive since it does not participate in OSPF advertisements.

Router 1 and 2 must enable link-detect, to stop advertising 1.1.1.0/24 in the event of a link failure.

22.1.5.2. VRRP Operation

Router 1 and 2 have VRRP setup on their Ethernet connection so that they can both function as the gateway for the clients on their network segment. Normally Router 1 is the VRRP master, and only in case of a link failure to the switch or the router failing, will Router 2 take over the virtual IP. The virtual IP used as the gateway is 1.1.2.254. Each router also has its own IP on the network so that each can be reached individually.

If Router 1 or its Ethernet link fail, VRRP will detect the link being down and remove the direct route to the 1.1.2.0/24. VRRP on Router 2 will stop seeing messages from Router 1, elect itself master and will take over the gateway for the network.

OSPF on router 1 will notice the link being down (and the route to 1.1.2.0/24 disappearing) and will use information from router 2 install a route to 1.1.2.0/24 via Router 2.

Router 3 will notice that Router 2 is now a more direct path to 1.1.2.0/24 network and start sending to Router 2 instead of Router 1.

After the failure all routers still know how to reach the entire network, and the clients on 1.1.2.0/24 can still send on the network using the same gateway address. The clients will see only a MAC address change of the gateway and experience a few seconds of network outage When the link returns, VRRP will switch back to the master, and the routes will return to their normal state.

Note that if the Router 1 WAN link fails, Router will see routes to Router3 via the Router 1 – Router 2 WAN and Ethernet links. If the faster Router 1 – Router 2 Ethernet path fails, Router 1 will fall back to the Router 1 – Router 2 WAN link.

Note that it would not be useful to leave the Ethernet 1.1.2.0/24 subnets out of the area and turn on redistribute connected as OSPF would not use the subnets for routing.

22.1.6. BGP Fundamentals

The Border Gateway Protocol (BGP, RFC 4271) is a robust and scalable routing protocol. BGP is designed to manage a routing table of up to 90000 routes. Therefore, it is used in large networks or among groups of networks which have common administrative and routing policies. If BGP is used to exchange routing information between different networks, it is called Exterior BGP (EBGP). Interior BGP (IBGP) is used to exchange routing information between routers within the same network.

22.2. Dynamic Routing Configuration

Figure 22.2. Dynamic Routing Menu

The Dynamic Routing menu is accessible from the main menu under *routing*. The path to this menu is *routing/dynamic*.

The *BGP*, *RIP* and *OSPF* menus provide access to features that enable configuration of the corresponding routing protocols.

22.3. RIP

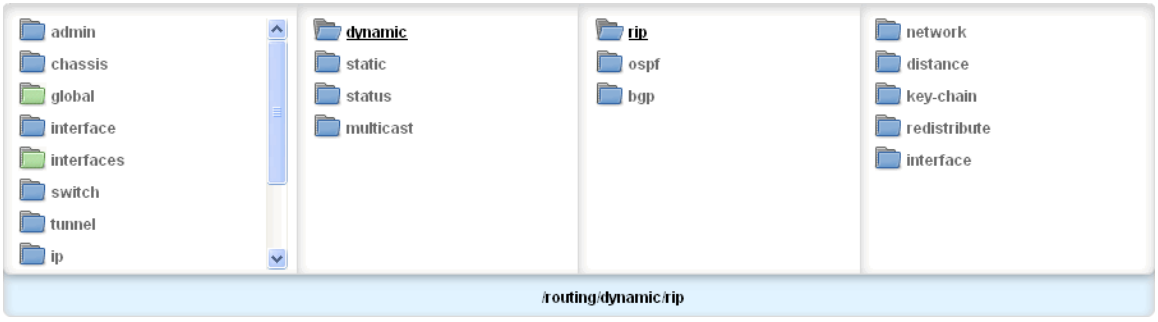

Figure 22.3. RIP Menu

The RIP menu is accessible from the main menu under *routing*. The path to this menu is *routing/dynamic/rip*.

22.3.1. RIP Configuration

The RIP Configuration form and Routing Timers form appear on the same screen as the *RIP* menu.

Figure 22.4. RIP Configuration Form

Enable RIP

Enables the RIP dynamic routing protocol.

Default Information Originate

The route element makes a static route only inside RIP. This element should be used only by advanced users who are particularly knowledgeable about the RIP protocol. In most cases, we recommend creating a static route and redistributing it in RIP using the redistribute element with static type.

Default Metric

Synopsis: *integer in the range [-32768 to 32767]*

Default: *1*

This element modifies the default metric value for redistributed routes. The value does not affect connected routes even if it is redistributed by redistribute connected. To modify the connected route's metric value, please use the redistribute connected metric.

Distance Default

Synopsis: *unsigned integer*

Set the default RIP distance to a specified value.

Version

Synopsis: *integer in the range [-32768 to 32767]*

Set the RIP version to accept for reads and send. The version can be either 1 or 2.

Disabling RIPv1 by specifying version 2 is STRONGLY encouraged.

Figure 22.5. Routing Timers Form

The RIP protocol has several timers. The user can configure those timers' values by the timers-basic element.

The default settings for the timers are as follows:

* The update timer is 30 seconds. Every update timer seconds, the RIP process is awakened to send an unsolicited response message containing the complete routing table to all neighboring RIP routers.

* The timeout timer is 180 seconds. Upon expiration of the timeout, the route is no longer valid; however, it is retained in the routing table for a short time so that neighbors can be notified that the route has been dropped.

* The garbage collect timer is 120 seconds. Upon expiration of the garbage-collection timer, the route is finally removed from the routing table.

Update Timer

Synopsis: *unsigned integer*

Default: *30*

The routing table update timer (in seconds).

Timeout Timer

Synopsis: *unsigned integer*

Default: *180*

The routing information timeout timer (in seconds).

Garbage Collection Timer

Synopsis: *unsigned integer*

Default: *120*

The garbage collection timer (in seconds).

22.3.1.1. RIP Networks

Neighbors are specific routers with which to exchange routes using the RIP protocol. This can be used when you want to explicitly control which routers are part of your RIP network.

Networks are used when you want to add any router that is part of a specific subnet, or connected to a specific network interface to be part of your RIP network.

Both neighbors and networks can be used at the same time.

For point to point links (T1/E1 links for example), one must use neighbor entries to add other routers to exchange routes with. Also note that RIP v1 does not send subnet mask information in its updates. Any defined networks are restricted to the classic (in the sense of Class A, B and C) networks. RIP v2 does not have this failing.

Subnet

Subnet Address/Prefix

Synopsis: *IPv4 address and prefix in CIDR notation*

Network address/prefix.

Neighbor

Neighbor IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

Neighbor IP address.

22.3.1.2. Distance

Distance with Matched Subnet

Subnet/Prefix

Synopsis: IPv4 address and prefix in CIDR notation

IP Address/Prefix.

Distance

Synopsis: unsigned integer

Distance value.

22.3.1.3. RIP Key Chains

The Key Chains table configures authentication keys used on the interfaces. By defining the keys in a key chain, the same settings can be applied to multiple groups of interfaces. Without key chains the same settings would have to be entered for each interface separately.

Key chains also allow multiple keys to be entered in a single key chain with a start time for when that key should become valid as well as the duration the key is valid. This allows multiple keys to be set up with automatic transitions from one key to the next over time.

A key consists of a key string, which is the value used for authentication. It also has the optional lifetime to accept RIP messages with the key, and the optional lifetime to send RIP messages with that key.

Key chain management.

Key Chain Name

Synopsis: *string*

The key chain name.

Key Configuration

Configure a key.

Key ID

Synopsis: *unsigned integer*

The key identifier number.

Key

Synopsis: *string*

Sets the key string.

Accept Lifetime

Set the accept lifetime of the key.

Time to Start

Synopsis: *date and time specification*

The time to start.

Expire Time

Synopsis: *date and time specification*

Synopsis: *string - the keyword { infinite }*

Expire time.

Send Lifetime

Set the send lifetime of the key.

Time to Start

Synopsis: *date and time specification*

The time to start.

Expire Time

Synopsis: *date and time specification*

Synopsis: *string - the keyword { infinite }*

Expire time.

22.3.1.4. Redistribute

This element redistributes routing information into the RIP tables from route entries specified by type.

Redistribute Route from other Protocols

Redistribute Type

Synopsis: *string - one of the following keywords { bgp, ospf, connected, static, kernel }*

Redistribute route type.

Metric

Synopsis: *integer in the range [-32768 to 32767]*

The metric for redistributed routes.

22.3.1.5. RIP Interfaces

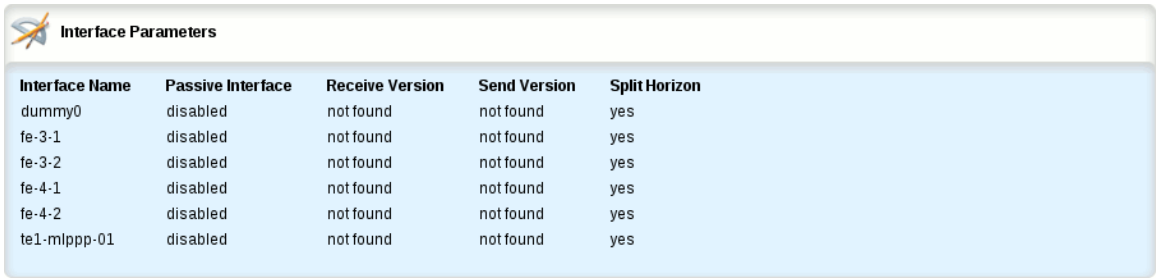

Interface Parameters

Interface Name	Passive Interface	Receive Version	Send Version	Split Horizon
dummy0	disabled	not found	not found	yes
fe-3-1	disabled	not found	not found	yes
fe-3-2	disabled	not found	not found	yes
fe-4-1	disabled	not found	not found	yes
fe-4-2	disabled	not found	not found	yes
te1-mlppp-01	disabled	not found	not found	yes

Figure 22.6. RIP Interface Parameters Table

RIP parameters refer to RIP parameters on the selected interface. The path to the RIP Interface Parameters table is *routing/dynamic/rip/interface*.

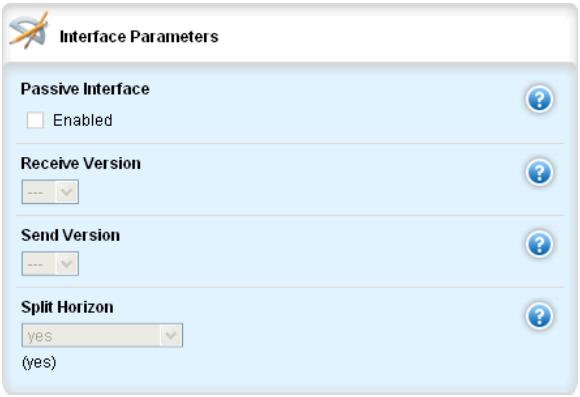

Figure 22.7. RIP Interface Parameters Form

To display the RIP Interface Parameters form and the Authentication form, navigate to *routing/dynamic/rip/interface/{interface}*.

Passive Interface

The specified interface is set to passive mode. In passive mode, all received packets are processed normally and ripd sends neither multicast nor unicast RIP packets except to RIP neighbors specified with a neighbor element.

Receive Version

Synopsis: *string - one of the following keywords { 2,1, 1,2, 2, 1 }*

The version of RIP packets that will be accepted on this interface. By default, version 1 and version 2 packet will be accepted.

Send Version

Synopsis: *string - one of the following keywords { 2,1, 1,2, 2, 1 }*

The version of RIP to send packets with. By default, version 2 packet will be sent.

Split Horizon

Synopsis: *string - one of the following keywords { poisoned-reverse, no, yes }*

Default: *yes*

A split horizon.

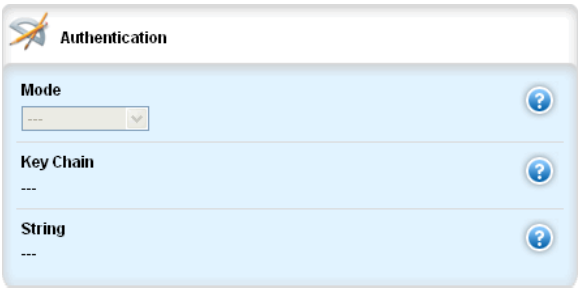

Figure 22.8. Authentication Form

Mode

Synopsis: *string - one of the following keywords { none, text, md5-old-ripd, md5-rfc }*

The authentication mode.

Key Chain

Synopsis: *string*

The authentication key-chain.

String

Synopsis: *string*

The authentication string.

> *Split horizon controls whether routes learned through an interface should be allowed to be advertised back out that interface. By default RIP advertises all routes it knows about to everyone, which makes it take a very long time for dropped links to age out of the network. The split horizon prevents advertising those routes back out the same interface which helps to control this problem. Some network topologies with rings of routers will still have some issues with aging out dead routes even with split horizon enabled but they will still age out faster. If fast network recovery is desired, use OSPF.*

22.4. OSPF

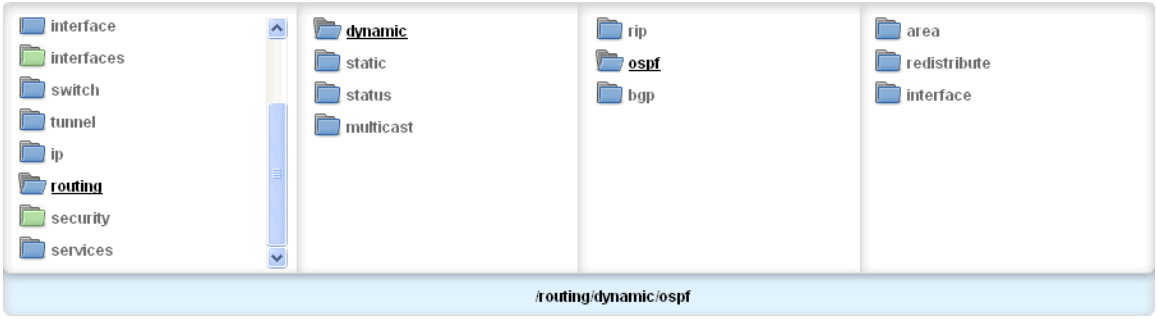

Figure 22.9. OSPF Menu

The OSPF menu is accessible from the main menu under *routing*. The path to this menu is *routing/dynamic/ospf*. The OSPF Area Distance form and the OSPF Configuration form appear on the same screen as the OSPF menu.

22.4.1. OSPF Configuration

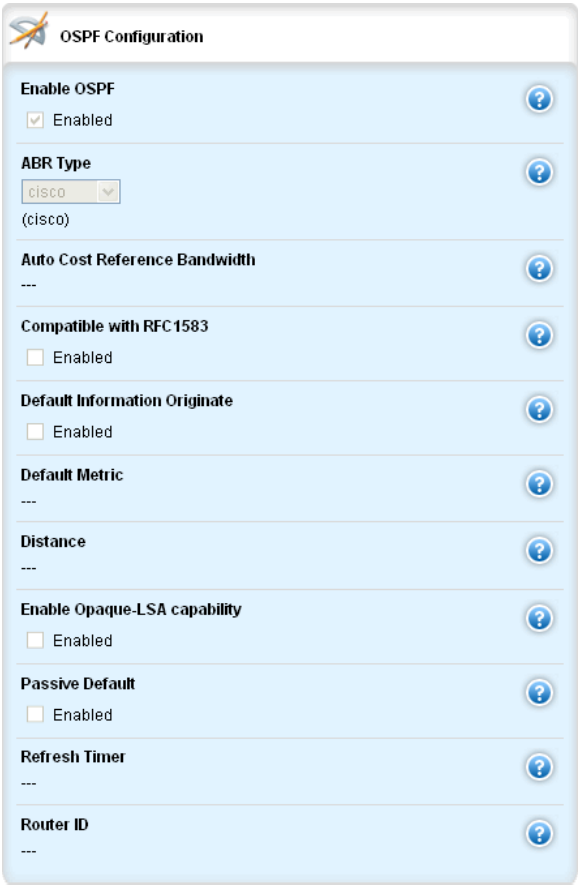

Figure 22.10. OSPF Configuration Form

Enable OSPF

Enables the OSPF dynamic routing protocol.

ABR Type

Synopsis: *string - one of the following keywords { standard, shortcut, ibm, cisco }*

Default: *cisco*

The OSPF ABR type.

Auto Cost Reference Bandwidth

Synopsis: *unsigned integer*

Default: *100*

Calculates the OSPF interface cost according to bandwidth [1-4294967 Mbps]

Compatible with RFC1583

Enables the compatibility with the obsolete RFC1583 OSPF (the current is RFC2178)

Default Information Originate

Advertises the default route.

Default Metric

Synopsis: *unsigned integer*

The default metric of redistribute routes.

Distance

Synopsis: *unsigned integer*

The administrative distance.

Enable Opaque-LSA capability

Enables the Opaque-LSA capability (rfc2370).

Passive Default

Suppresses routing updates on interfaces by default.

Refresh Timer

Synopsis: *unsigned short integer*

Default: *10*

The refresh timer.

Router ID

Synopsis: *IPv4 address in dotted-decimal notation*

The Router ID for OSPF.

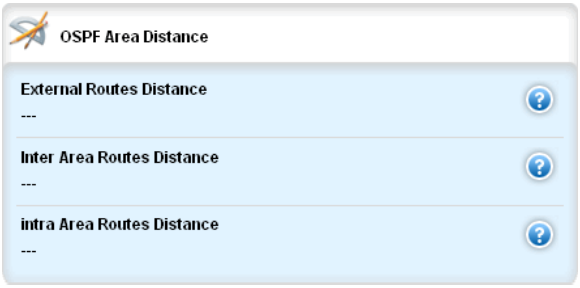

Figure 22.11. OSPF Area Distance Form

The OSPF Area Distance form can be used to define OSPF external, inter-area or intra-area routes distance.

External Routes Distance

Synopsis: *unsigned integer*

The administrative distance for external routes.

Inter Area Routes Distance

Synopsis: *unsigned integer*

The administrative distance for inter-area routes.

intra Area Routes Distance

Synopsis: *unsigned integer*

The administrative distance for intra-area routes.

22.4.1.1. OSPF Network Areas

OSPF uses areas to control which routes are distributed between routers.

Area ID.

Synopsis: *IPv4 address in dotted-decimal notation*

The OSPF Area ID (format: A.B.C.D).

Area Network/Prefix

Synopsis: *IPv4 address and prefix in CIDR notation*

The OSPF area network/prefix.

22.4.1.2. OSPF Redistribute

Redistribute from other Routing Protocols

This feature redistributes information from another routing protocol.

Redistribute Route From

Synopsis: *string - one of the following keywords { bgp, rip, connected, static, kernel }*

Redistributes the route type.

Metric Type

Synopsis: *integer in the range [-32768 to 32767]*

Default: *2*

The OSPF exterior metric type for redistributed routes.

Metric

Synopsis: *unsigned integer*

The metric for redistributed routes.

22.4.1.3. OSPF Interfaces

Figure 22.12. Interface Parameters Table

The path to the Interface Parameters table is *routing/dynamic/ospf/interface*.

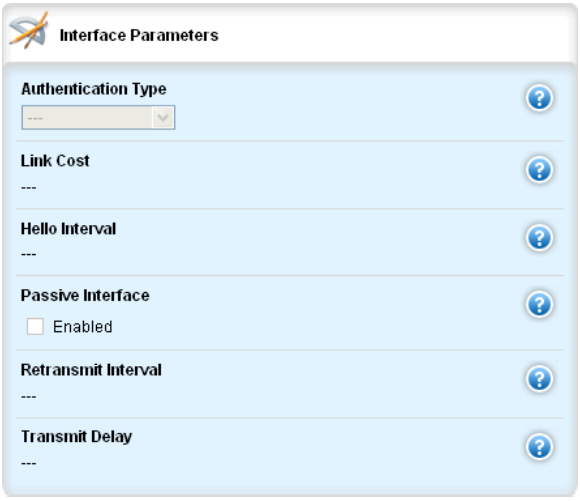

Figure 22.13. Interface Parameters Form

The path to the Interfaces form and Dead Interval form is *routing/dynamic/ospf/interface/{interface}*.

OSPF parameters on the selected interface.

Interface Name

Synopsis: *string*

Interface name.

Authentication Type

Synopsis: *string - one of the following keywords { null, message-digest }*

The authentication type on this interface.

Link Cost

Synopsis: *unsigned integer*

The link cost. If not set, it cost is based on reference bandwidth.

Hello Interval

Synopsis: *unsigned integer*

Default: *10*

The time (in seconds) between sending hello packets.

priority

Synopsis: *unsigned byte integer*

Default: *1*

Priority of interface.

Passive Interface

Whether an interface is active or passive. Passive interfaces do not send LSAs to other routers and are not part of an OSPF area.

Retransmit Interval

Synopsis: *unsigned integer*

Default: *5*

Time (in seconds) between retransmitting lost link state advertisements.

Transmit Delay

Synopsis: *unsigned integer*

Default: *1*

The link state transmit delay (in seconds).

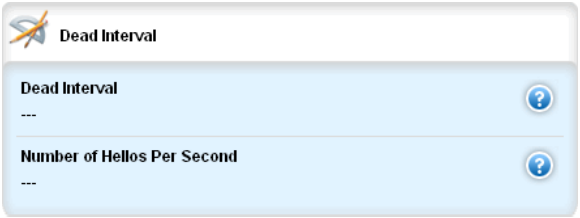

Figure 22.14. Dead Interval Form

A dead interval is the time before considering a router dead.

Dead Interval

Synopsis: *unsigned integer*

Default: *40*

The time before considering a router dead (in seconds).

Number of Hellos Per Second

Synopsis: *unsigned byte integer*

The number of times a hello message can be sent within one second.

Configuration Parameters

Configuration forms and display tables can be found at *routing/dynamic/ospf/interface*, then clicking on one of the interface submenus (for example, dummy0) and then clicking on the further set of submenus that follow (authentication-ip, cost-ip, dead-interval-ip, hello-interval-ip, message-digest-key, message-digest-key-ip, retransmit-interval-ip and transmit-delay-ip).

22.5. BGP

22.5.1. BGP configuration

Figure 22.15. BGP Menu

To display the BGP menu, navigate to *routing/dynamic/bgp*.

Figure 22.16. BGP Configuration Form

The BGP Configuration form appears on the same screen as the *BGP* menu.

Enable BGP

Enables BGP.

Autonomous System ID

Synopsis: *unsigned integer*

Autonomous System ID.

Always compare MED

Always comparing MED from different neighbors.

Default Local Preference

Synopsis: *unsigned integer*

Default: *100*

Default local preference value.

Deterministic Med

Pick the best-MED path among paths advertised from neighboring AS.

Router ID

Synopsis: *IPv4 address in dotted-decimal notation*

Router ID.

Figure 22.17. Distance Form

The path to the Distance form is *routing/dynamic/bgp*. Distance represents the distance value of BGP.

External Routes Distance

Synopsis: *unsigned integer*

Distance value for external routes.

Internal Routes Distance

Synopsis: *unsigned integer*

Distance value for internal routes.

Local Routes Distance

Synopsis: *unsigned integer*

Distance value for local routes.

22.5.1.1. Filter

The following information is available in configuration forms and display tables under the Filter submenu.

Autonomous System Path Filter

Name

Synopsis: *A string conforming to: "[^\s]+"*

Name of the AS-path filter.

Autonomous System Path Entry

Action

Synopsis: *string - one of the following keywords { permit, deny }*

Action.

Match

Synopsis: *string*

The regular expression to match the BGP AS paths.

Prefix List

Prefix List

Synopsis: *A string conforming to: "[^\s]+"*

The name of the prefix list.

Description

Synopsis: *string*

The description of the prefix list.

Prefix List Entry

Sequence Number

Synopsis: *unsigned integer*

Sequence number of the entry.

Action

Synopsis: *string - one of the following keywords { permit, deny }*

Default: *permit*

Action.

Network

Synopsis: *IPv4 address and prefix in CIDR notation*

Network (xxx.xxx.xxx.xxx/xx).

Less Than or Equal to

Synopsis: *unsigned byte integer*

The maximum prefix length to be matched.

Greater Than or Equal to

Synopsis: *unsigned byte integer*

The minimum prefix length to be matched.

Route Map

Route Map Tag

Synopsis: *A string conforming to: "[^\s]+"*

Route map tag.

Route Map Entry

Sequence Number

Synopsis: *unsigned integer*

The sequence number of the route-map entry.

Action

Synopsis: *string - one of the following keywords { permit, deny }*

Default: *permit*

Action.

Call Route Map

Synopsis: *A string conforming to: "[^\s]+"*

Jump to another route-map after match+set.

On Match Goto

Synopsis: *unsigned integer*

Go to this entry on match.

Route Map Match

AS Path Filter

Synopsis: *A string conforming to: "[^\s]+"*

Match the BGP AS path filter.

Match Address of Route

Prefix List

Synopsis: *A string conforming to: "[^\s]+"*

The prefix list name.

Match Next-Hop of Route

Prefix List

Synopsis: *A string conforming to: "[^\s]+"*

The prefix list name.

Route Source Match

Prefix List

Synopsis: *A string conforming to: "[^\s]+"*

The prefix list name.

Route Map Metric

Metric

Synopsis: *unsigned integer*

Match the route metric.

Peer Address

Synopsis: *IPv4 address in dotted-decimal notation*

Match the peer address.

Origin

Synopsis: *string - one of the following keywords { incomplete, igp, egp }*

Match the BGP origin code.

Aggregator

AS Number

Synopsis: *unsigned integer*

AS number.

IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

IP address of aggregator.

Prepend

AS Number

Synopsis: *unsigned integer*

AS number.

Exclude

AS Number

Synopsis: *unsigned integer*

AS number.

Local Preference

Synopsis: *unsigned integer*

Local preference.

Metric

operation

Synopsis: *string - one of the following keywords { sub, add, set }*

Set , add or subtract the metric value.

value

Synopsis: *unsigned integer*

Value.

next-hop

Synopsis: *IPv4 address in dotted-decimal notation*

Synopsis: *string - the keyword { peer }*

The next hop address (xxx.xxx.xxx.xxx/xx or peer to use peer address).

origin

Synopsis: *string - one of the following keywords { incomplete, igp, egp }*

The origin code.

originator-id

Synopsis: *IPv4 address in dotted-decimal notation*

The originator ID.

weight

Synopsis: *unsigned integer*

Weight.

22.5.1.2. Network

Networks may be specified in order to add BGP routers connected to the specified subnets. Note that a network specification need not be a given valid entry in the routing table. Since BGP is a border gateway protocol, one would more typically enter a more general network specification here. For example, if a routed network inside the AS comprised many different Class C subnets (/24) of the 192.168.0.0/16 range, it would be more efficient to advertise the one Class B network specification, 192.168.0.0/16, to one's BGP neighbors.

If BGP Neighbors are specified but no Networks are specified, then the router will receive BGP routing information from its neighbors but will not advertise any routes to them.

Subnet Address/Prefix

Synopsis: *IPv4 address and prefix in CIDR notation*

IP address/prefix.

22.5.1.3. Neighbor

Neighbors are other BGP routers with which to exchange routing information. One or more neighbors must be specified in order for BGP to operate.

If BGP Neighbors are specified but no Networks are specified, then the router will receive BGP routing information from its neighbors but will not advertise any routes to them.

Neighbor IP address

Synopsis: *IPv4 address in dotted-decimal notation*

The neighbor IP address.

Neighbor Autonomous System ID

Synopsis: *unsigned integer*

A BGP neighbor.

ebgp-multihop

Synopsis: *unsigned byte integer*

The maximum hop count. This allows EBGP neighbors not on directly connected networks.

Maximum Prefix

Synopsis: *unsigned integer*

The maximum prefix number accepted from this peer.

next-hop-self

Disables the next hop calculation for this neighbor.

password

Synopsis: *A string conforming to: "[^\s]+"*

Password.

soft-reconfiguration

Per neighbor soft reconfiguration.

weight

Synopsis: *unsigned short integer*

The default weight for routes from this neighbor.

Route Map

in

Synopsis: *A string conforming to: "[^\s]+"*

Apply route map to incoming routes

out

Synopsis: *A string conforming to: "[^\s]+"*

Apply route map to outbound routes.

22.5.1.4. Aggregate-address

Aggregate Network

subnet

Synopsis: *IPv4 address and prefix in CIDR notation*

subnet (xxx.xxx.xxx.xxx/xx).

Options

value

Synopsis: *string - one of the following keywords { summary-only, as-set }*

Aggregate address option.

22.5.1.5. Distance-ip

Distance with Matched Subnet

Subnet/Prefix

Synopsis: IPv4 address and prefix in CIDR notation

IP Address/Prefix.

Distance

Synopsis: unsigned integer

Distance value.

22.5.1.6. Redistribute

Redistribute Route from Other Protocols

Redistribute Route From

Synopsis: *string - one of the following keywords { rip, ospf, connected, static, kernel }*

Redistribute route type.

Metric

Synopsis: *unsigned integer*

Metric value for redistributed routes.

23. Static Routing

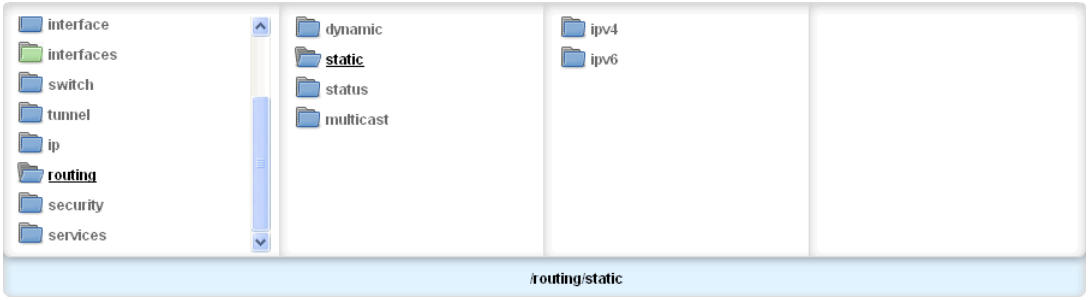

Figure 23.1. Static Menu

The Static menu is accessible from the main menu under *routing*. The path to this menu is *routing/static*.

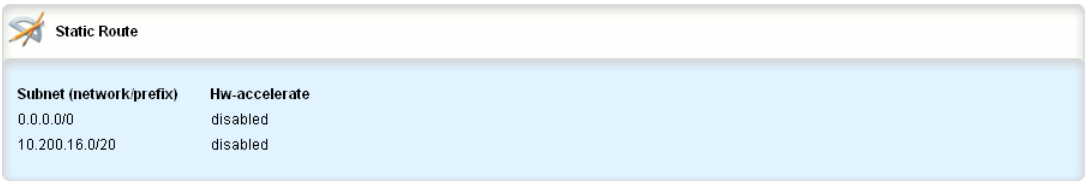

Figure 23.2. Static Route table

The path to the Static Route table is *routing/static/ipv4*.

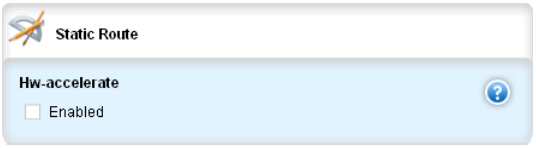

Figure 23.3. Static Route form

The path to the Static Route form is *routing/static/ipv4/{route}*.

hw-accelerate

If the static unicast route can be hardware accelerated, the option will be available. For a static unicast route to be accelerated, the ingress and egress interfaces must be switched.

The Hw-accelerate field will only be visible in the form if the device has a layer 3 switch.

Figure 23.4. Static Route Using Gateway table

The path to the Static Route Using Gateway table is *routing/static/ipv4/{route}/via*.

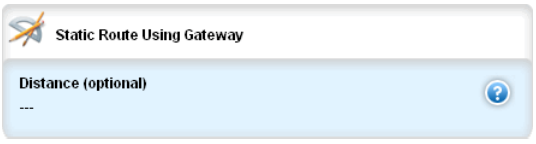

Figure 23.5. Static Route Using Gateway form

The path to the Static Route Using Gateway form is *routing/static/ipv4/{route}/via/{address}*.

Gateway Address

Synopsis: *IPv4 address in dotted-decimal notation*

The gateway for the static route.

Distance (optional)

Synopsis: *unsigned integer*

The distance for the static route.

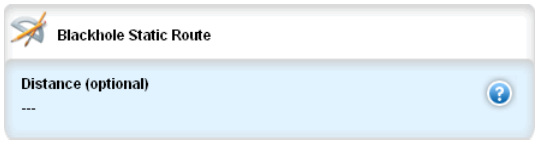

Figure 23.6. Blackhole Static Route form

The path to the Blackhole Static Route form is *routing/static/ipv4/{route}/blackhole*.

Distance (optional)

Synopsis: *unsigned integer*

The distance for the static route.

Figure 23.7. Static Route Using Interface table

The path to the Static Route Using Interface table is *routing/static/ipv4/{route}/dev*.

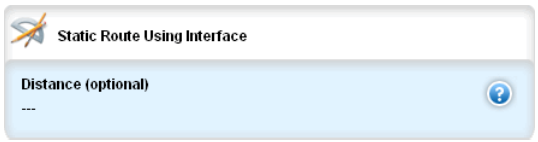

Figure 23.8. Static Route Using Interface form

The path to the Static Route Using Interface form is *routing/static/ipv4/{route}/dev/{interface}*.

Interface Name

Synopsis: *A string*

The interface for the static route.

Distance (optional)

Synopsis: *unsigned integer*

The distance for the static route.

24. Routing Status

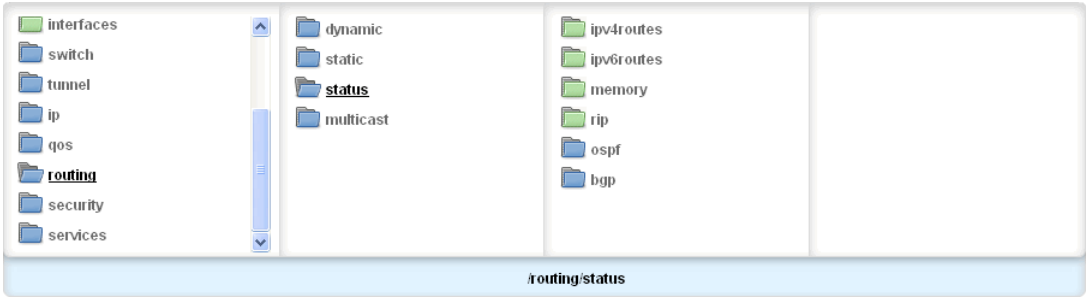

Figure 24.1. Routing Status Menu

The Routing Status menu is accessible under *routing/status*.

24.1. IPv4

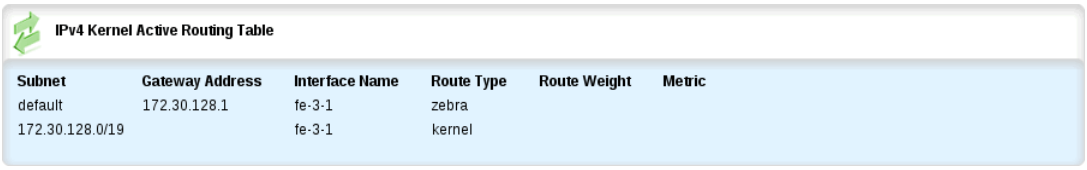

Figure 24.2. IPv4 Kernel Active Routing Table

The path to the IPv4 Kernel Active Routing table is *routing/status/ipv4routes*.

> *It is possible to create a route on a locally connected broadcast network (i.e. without a gateway) without also bringing up a corresponding IP address on that interface. For example, it would be possible to add* ***192.168.1.0/24*** *to* ***fe-3-1****, which has an IP address of* ***10.0.1.1*** *but no corresponding alias address on the* ***192.168.1.0/24*** *subnet.*

Subnet

Synopsis: *string*

The network/prefix.

Gateway Address

Synopsis: *string*

The gateway address.

Interface Name

Synopsis: *string*

The interface name.

Route Type

Synopsis: *string*

The route type.

Route Weight

Synopsis: *string*

The route weight.

Metric

Synopsis: *string*

The route metric value.

24.2. IPv6

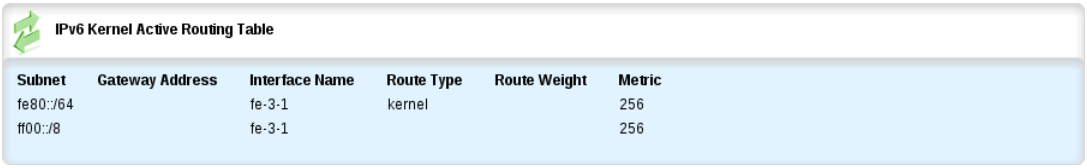

Figure 24.3. IPv6Kernel Active Routing Table

The path to the IPv6 Kernel Active Routing table is *routing/status/ipv6routes*.

Subnet

Synopsis: *string*

The network/prefix.

Gateway Address

Synopsis: *string*

The gateway address.

Interface Name

Synopsis: *string*

The interface name.

Route Type

Synopsis: *string*

The route type.

Route Weight

Synopsis: *string*

The route weight.

Metric

Synopsis: *string*

The metric value.

24.3. Memory Statistics

The path to the memory forms is *routing/status/memory*.

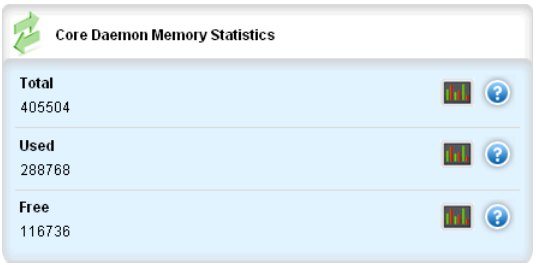

Figure 24.4. Core Daemon Memory Statistics Form

Total heap allocated (Byte)

Synopsis: *unsigned integer*

The total heap allocated (in bytes).

Used ordinary blocks (Byte)

Synopsis: *unsigned integer*

The number of used ordinary blocks (in bytes).

Free ordinary blocks (Byte)

Synopsis: *unsigned integer*

The number of free ordinary blocks (in bytes).

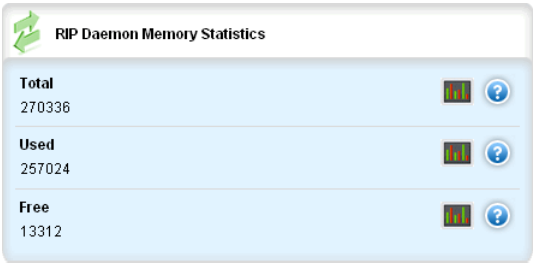

Figure 24.5. RIP Daemon Memory Statistics Form

total

Synopsis: *unsigned integer*

The total heap allocated (in bytes).

used

Synopsis: *unsigned integer*

The number of used ordinary blocks (in bytes).

free

Synopsis: *unsigned integer*

The number of free ordinary blocks (in bytes).

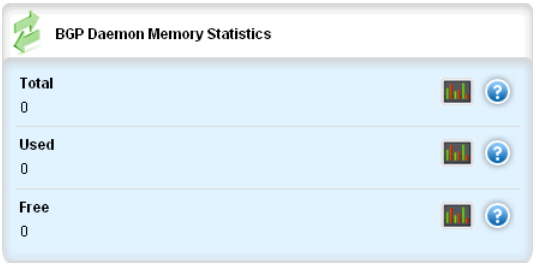

Figure 24.6. BGP Daemon Memory Statistics Form

total

Synopsis: *unsigned integer*

The total heap allocated (in bytes).

used

Synopsis: *unsigned integer*

The number of used ordinary blocks (in bytes).

free

Synopsis: *unsigned integer*

The number of free ordinary blocks (in bytes).

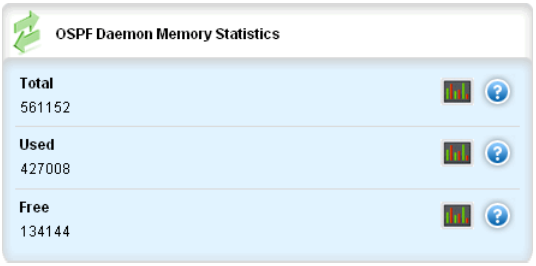

Figure 24.7. OSPF Daemon Memory Statistics Form

total

Synopsis: *unsigned integer*

The total heap allocated (in bytes).

used

Synopsis: *unsigned integer*

The number of used ordinary blocks (in bytes).

free

Synopsis: *unsigned integer*

The number of free ordinary blocks (in bytes).

24.4. RIP

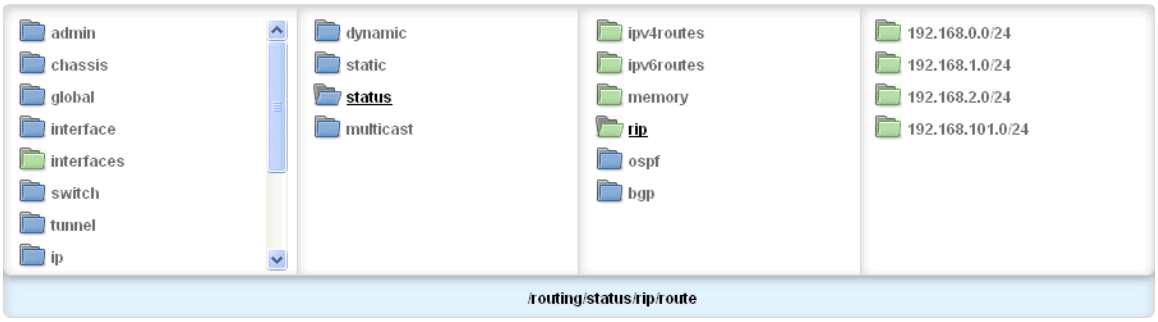

Figure 24.8. RIP Menu

The RIP status tables and forms are accessible from the RIP menu at *routing/status/rip* and *routing/status/rip/{address}*.

24.5. OSPF

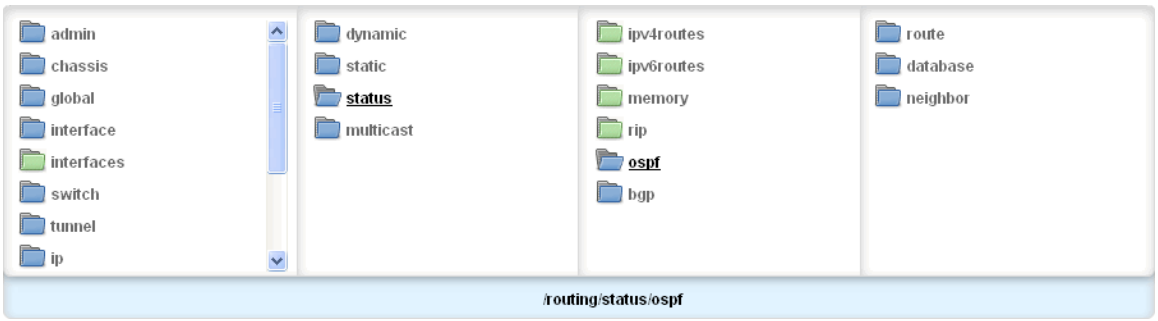

Figure 24.9. OSPF Menu

To display the OSPF menu, navigate to *routing/status/ospf*.

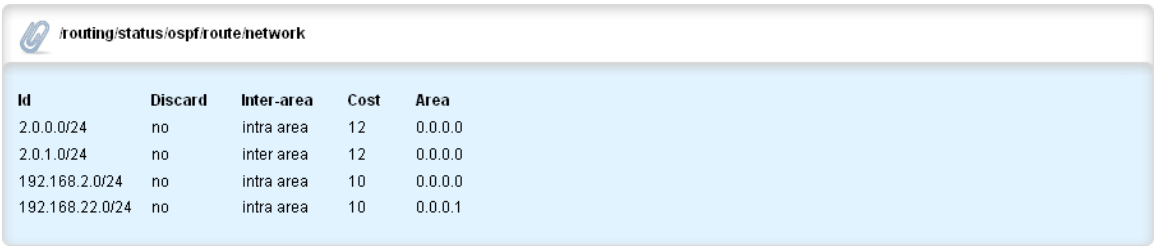
/routing/status/ospf/route/network

Id	Discard	Inter-area	Cost	Area
2.0.0.0/24	no	intra area	12	0.0.0.0
2.0.1.0/24	no	inter area	12	0.0.0.0
192.168.2.0/24	no	intra area	10	0.0.0.0
192.168.22.0/24	no	intra area	10	0.0.0.1

Figure 24.10. Network Table

To display the Network table, navigate to *routing/status/ospf/route/network*.

id

Synopsis: *string*

Network Prefix.

discard

Synopsis: *string*

This entry is discarded entry.

inter-area

Synopsis: *string*

Is path type inter area.

cost

Synopsis: *string*

Cost.

area

Synopsis: *string*

Area.

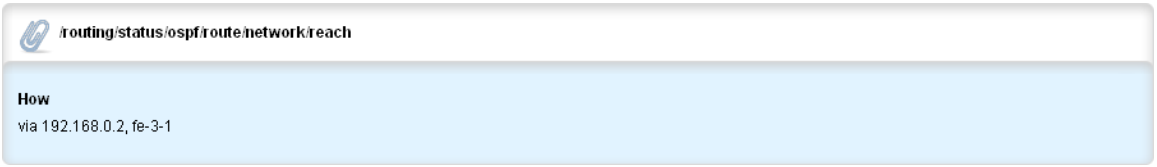
/routing/status/ospf/route/network/reach

How
via 192.168.0.2, fe-3-1

Figure 24.11. Reach Table

To display the Reach table, navigate to *routing/status/ospf/route/network/{address}/reach*.

how

Synopsis: *string*

How to reach this network.

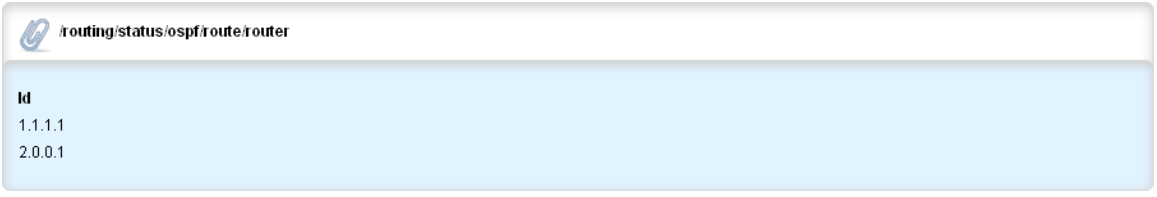
/routing/status/ospf/route/router

Id
1.1.1.1
2.0.0.1

Figure 24.12. Router Table

To display the Router table, navigate to *routing/status/ospf/route/router*.

id

Synopsis: *string*

Router ID.

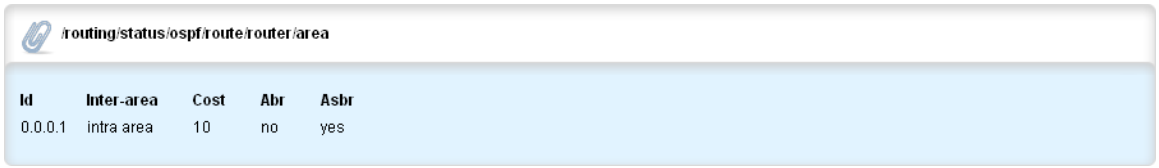

Figure 24.13. Area Table

To display the Area table, navigate to *routing/status/ospf/route/router/{number}/area*.

id

Synopsis: *string*

Area ID.

inter-area

Synopsis: *string*

Is path type inter area.

cost

Synopsis: *string*

Cost.

abr

Synopsis: *string*

Is area border router.

asbr

Synopsis: *string*

Is autonomous system border router.

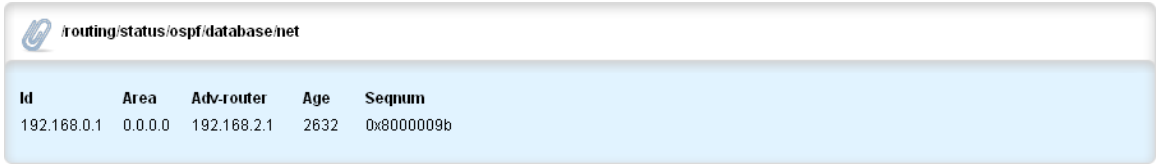

Figure 24.14. Net Table

To display the Net table, navigate to *routing/status/ospf/database/net*.

id

Synopsis: *string*

Link ID.

area

Synopsis: *string*

Area ID.

adv-router

Synopsis: *string*

Advertising Router.

age

Synopsis: *integer*

Age.

seqnum

Synopsis: *string*

Sequence number.

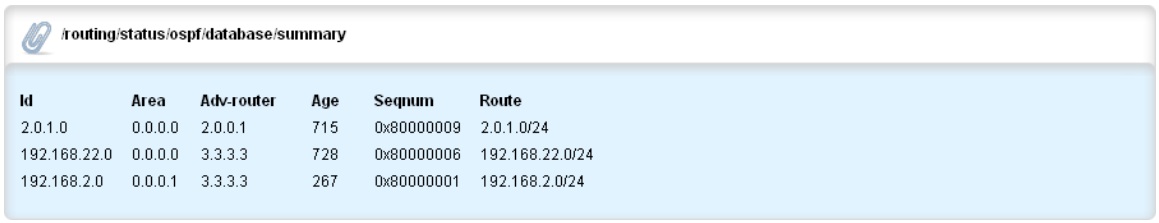

Figure 24.15. Summary Table

To display the Summary table, navigate to *routing/status/ospf/database/summary*.

id

Synopsis: *string*

Link ID.

area

Synopsis: *string*

Area ID.

adv-router

Synopsis: *string*

Advertising Router.

age

Synopsis: *integer*

Age.

seqnum

Synopsis: *string*

Sequence number.

route

Synopsis: *string*

Route.

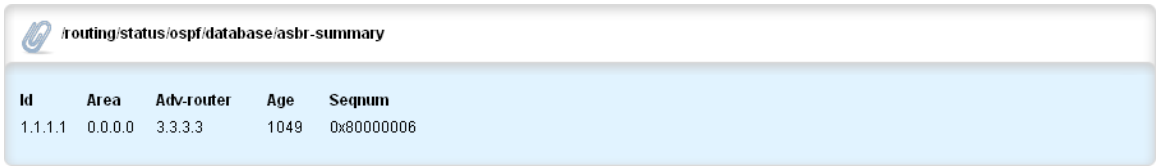

Figure 24.16. ASBR-Summary Table

To display the ASBR-Summary table, navigate to *routing/status/ospf/asbr-summary*.

id

Synopsis: *string*

Link ID.

area

Synopsis: *string*

Area ID.

adv-router

Synopsis: *string*

Advertising Router.

age

Synopsis: *integer*

Age.

seqnum

Synopsis: *string*

Sequence number.

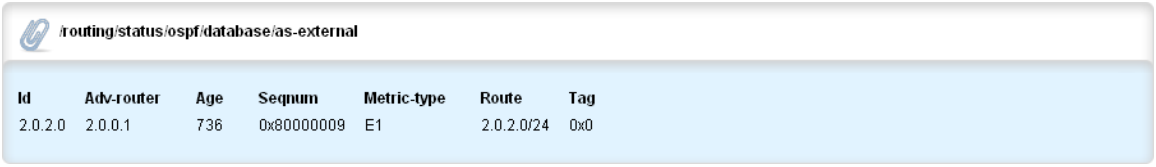
/routing/status/ospf/database/as-external
Id Adv-router Age Seqnum Metric-type Route Tag
2.0.2.0 2.0.0.1 736 0x80000009 E1 2.0.2.0/24 0x0

Figure 24.17. AS-External Table

To display the AS-External table, navigate to *routing/status/ospf/database/as-external*.

id

Synopsis: *string*

Link ID.

adv-router

Synopsis: *string*

Advertising Router.

age

Synopsis: *integer*

Age.

seqnum

Synopsis: *string*

Sequence number.

metric-type

Synopsis: *string*

External metric type.

route

Synopsis: *string*

Route.

tag

Synopsis: *string*

Route tag.

/routing/status/ospf/neighbor
Id Address Priority State Dead-time Interface
1.1.1.1 192.168.22.1 1 Full/Backup 37.592s switch.0022:192.168.22.2

Figure 24.18. Neighbor Table

To display the Neighbor table, navigate to *routing/status/ospf/neighbor*.

id

Synopsis: *string*

Neighbor ID.

address

Synopsis: *string*

Address.

priority

Synopsis: *integer*

Priority.

state

Synopsis: *string*

State.

dead-time

Synopsis: *string*

Dead Time.

interface

Synopsis: *string*

Interface.

24.6. BGP

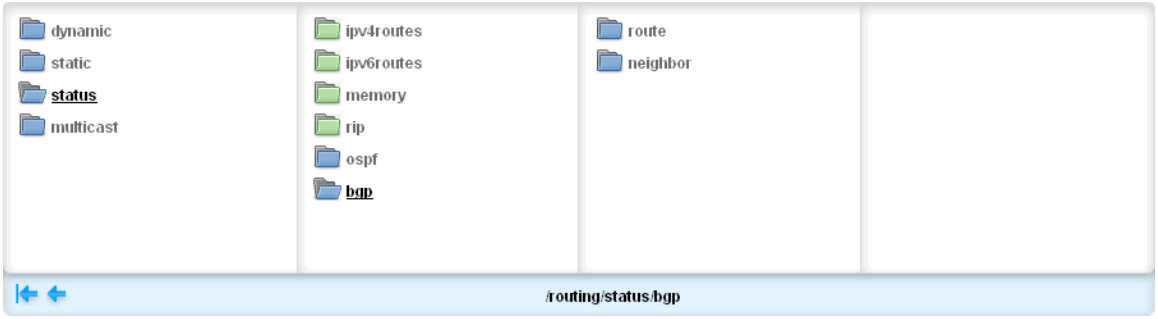

Figure 24.19. BGP Menu

To display the BGP menu, navigate to *routing/status/bgp*.

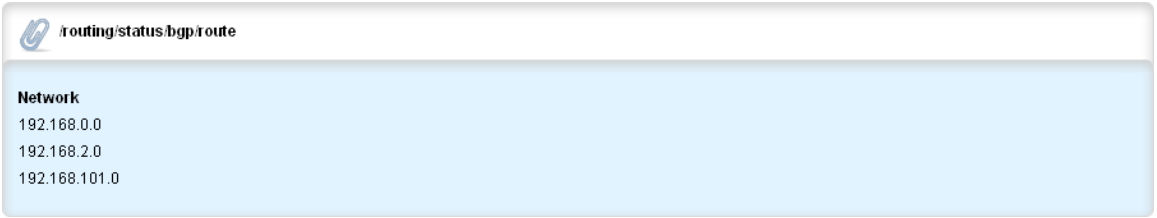

Figure 24.20. Route Table

To display the BGP Route table, navigate to *routing/status/bgp/route*.

network

Synopsis: *string*

Network.

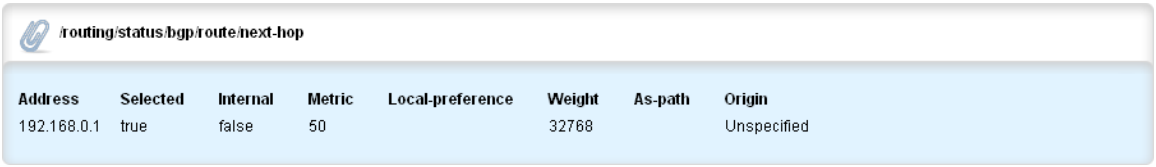
/routing/status/bgp/route/next-hop
Address Selected Internal Metric Local-preference Weight As-path Origin
192.168.0.1 true false 50 32768 Unspecified

Figure 24.21. Next Hop Table

To display the Next Hop table, navigate to *routing/status/bgp/route/{address}/next-hop*.

address

Synopsis: *string*

Next-hop address.

selected

Synopsis: *boolean*

Selected next-hop for this route.

internal

Synopsis: *boolean*

Internal route.

metric

Synopsis: *string*

Metric.

local-preference

Synopsis: *string*

Local preference.

weight

Synopsis: *integer*

Weight.

as-path

Synopsis: *string*

AS path.

origin

Synopsis: *string*

Origin.

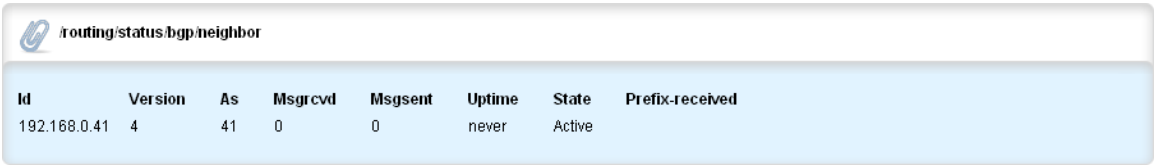

Figure 24.22. BGP Neighbor Table

To display the Neighbor table, navigate to *routing/status/bgp/neighbor*.

id

Synopsis: *string*

Neighbor address.

version

Synopsis: *integer*

BGP version.

as

Synopsis: *string*

Remote AS number.

msgrcvd

Synopsis: *integer*

Number of received BGP messages.

msgsent

Synopsis: *integer*

Number of sent BGP messages.

uptime

Synopsis: *string*

Peer up time.

state

Synopsis: *string*

Connection state with this neighbor.

prefix-received

Synopsis: *string*

Number of prefixes (networks) received from this neighbor.

25. Multicast Routing

Figure 25.1. Multicast Routing menu

The Multicast Routing menu is accessible from the main menu under *routing*. The path to this menu is *routing/multicast*. The user can choose between enabling dynamic multicast routing or static multicast routing by checking off "Enable" on the Routing Multicast Dynamic Form or the Routing Multicast Static Form.

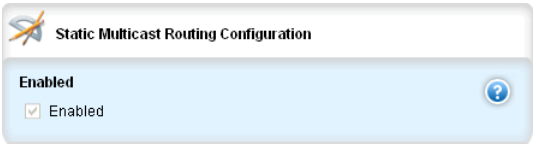

Figure 25.2. Static Multicast Routing Configuration form

The Static Multicast Routing Configuration form appears on the same screen as the *Multicast* menu.

enabled

Enables static multicast routing service

Figure 25.3. Static menu

The path to the Static menu is *routing/multicast/static*. From the *static* menu, there are two branches of menus. By clicking on the *mcast-groups* menu, the user can access forms used to create multicast groups for forwarding. Clicking on the *status* menu allows you to view the status of your configured multicast groups. The following forms and tables are linked to the *mcast-groups* menu. The features available from the *status* menu are covered a little later in this chapter.

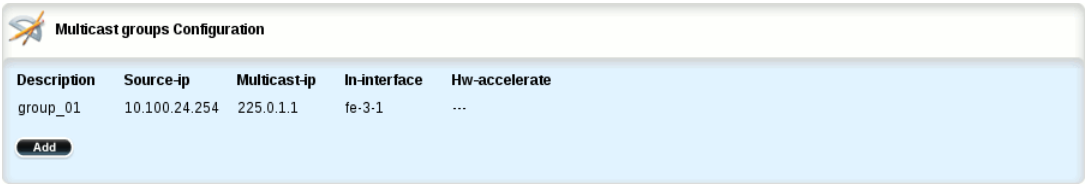

Figure 25.4. Multicast Groups Configuration table

The path to the Multicast Groups Configuration table is *routing/multicast/static/mcast-groups*.

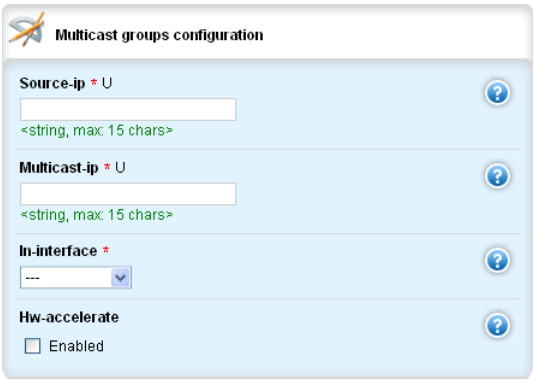

Figure 25.5. Multicast Groups Configuration form

The path to the Multicast Groups Configuration form is *routing/multicast/static/mcast-groups* and then clicking on one of the linked submenus.

description

Synopsis: *A string*

Describes this multicast group

source-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The expected source IP address of the multicast packet, in the format xxx.xxx.xxx.xxx

“U” indicates that this address is uniquely paired with the multicast address set in the **Multicast-ip** field. You cannot use this IP address to create another Multicast Routing entry with a different Multicast-ip address.

multicast-ip

Synopsis: *A string conforming to: "(22[4-9]|23[0-9])\.(([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])\.){2}([0-9]|[1-9][0-9]|1[0-9]{2}|2[0-4][0-9]|25[0-5])"*

The multicast IP address to be forwarded, in the format xxx.xxx.xxx.xxx

The address must be in the range of 224.0.0.0 to 239.255.255.255

“U” indicates that this address is uniquely paired with the source IP address set in the **Source-ip** field. You cannot use this IP address to create another Multicast Routing entry with a different Source-ip address.

in-interface

Synopsis: *A string*

The interface upon which the multicast packet arrives

hw-accelerate

If the multicast route can be hardware accelerated, the option will be available. For a multicast route to be accelerated, the ingress and egress interfaces must be switched.

Figure 25.6. Outgoing Interfaces table

The path to the Outgoing Interfaces table is *routing/multicast/static/mcast-groups* and then clicking on one of the linked submenus.

The OutInterface is the interface to which the matched multicast packet will be forwarded.

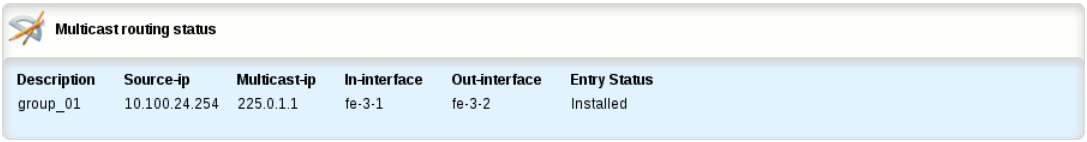

Figure 25.7. Multicast Routing Status table

The path to the Multicast Routing Status table is *routing/multicast/static/status*.

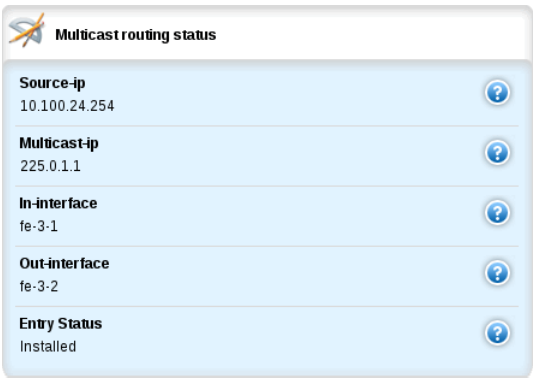

Figure 25.8. Multicast Routing Status form

The path to the Multicast Routing Status form is *routing/multicast/static/status* and then clicking on one of the linked submenus.

description

Synopsis: *A string*

Describes this multicast group

source-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The expected source IP address of the multicast packet, in the format xxx.xxx.xxx.xxx

“U” indicates that this address is uniquely paired with the multicast address set in the **Multicast-ip** field. You cannot use this IP address to create another Multicast Routing entry with a different Multicast-ip address.

multicast-ip

Synopsis: *IPv4 address in dotted-decimal notation*

The multicast IP address to be forwarded, in the format xxx.xxx.xxx.xxx

The address must be in the range of 224.0.0.0 to 239.255.255.255

“U” indicates that this address is uniquely paired with the source IP address set in the **Source-ip** field. You cannot use this IP address to create another Multicast Routing entry with a different Source-ip address.

in-interface

Synopsis: *A string*

The interface upon which the multicast packet arrives

out-interface

Synopsis: *string*

The interface to which the matched multicast packet will be forwarded

entryStatus

Synopsis: *string*

The status of the multicast routing entry

26. Firewall

26.1. Firewall Fundamentals

Firewalls are software systems designed to prevent unauthorized access to or from private networks. Firewalls are most often used to prevent unauthorized Internet users from accessing private networks (intranets) connected to the Internet.

When the ROX™ firewall is used, the router serves as a *gateway* machine through which all messages entering or leaving the intranet pass. The router examines each message and blocks those that do not meet the specified security criteria. The router also acts as a *proxy*, preventing direct communication between computers on the Internet and intranet. Proxy servers can filter the kinds of communication that are allowed between two computers and perform address translation.

26.1.1. Stateless vs Stateful Firewalls

Firewalls fall into two broad categories: *stateless* and *stateful* (session-based).

Stateless or “static” firewalls make decisions about traffic without regard to traffic history; they simply open a "hole" for the traffic’s type, based upon a TCP or UDP port number. Stateless firewalling is relatively simple, easily handling web and email traffic. However, stateless firewalls have some disadvantages. All holes opened in the firewall are always open, and connections are not opened or closed based on outside criteria. Static IP filters offer no form of authentication.

Stateful firewalling adds considerable complexity to the firewalling process by tracking the state of each connection.

A stateful firewall also looks at and tests each packet, and the tests or “rules” may be modified depending on packets that have already been processed. This is called “connection tracking”. Stateful firewalls can also recognize that traffic on connected sets of TCP/UDP ports is from a particular protocol and manage it as a whole.

26.1.2. Linux® netfilter, iptables, and the Firewall

ROX™ employs a stateful firewall system known as *netfilter*, a subsystem of the Linux kernel that provides the ability to examine IP packets on a per-session basis.

The netfilter system uses rulesets, which are collections of packet classification rules that determine the outcome of the examination of a specific packet. The rules are defined by *iptables*, a generic table structure syntax and utility program for the configuration and control of netfilter.

ROX™ implements an IP firewall using a structured user interface to configure iptables rules and netfilter rulesets.

26.1.3. Network Address Translation

Network Address Translation (NAT) enables a LAN to use one set of IP addresses for internal traffic and a second set for external traffic. The netfilter NAT function makes all necessary IP address translations as traffic passes between the intranet and Internet. NAT is often referred to in Linux as IP Masquerading.

NAT itself provides a type of firewall by hiding internal IP addresses. More importantly, NAT enables a network to use more internal IP addresses. Since they are used internally only, there is no possibility of conflict with IP addresses used by other organizations. Typically, an internal network is configured to use one or more of the reserved address blocks described in RFC1918:

IP Network/Mask	Address Range
10.0.0.0/8	10.0.0.0 - 10.255.255.255
172.16.0.0/12	172.16.0.0 - 172.31.255.255

IP Network/Mask	Address Range
192.168.0.0/16	192.168.0.0 - 192.168.255.255

Table 26.1. RFC1918 Reserved IP Address Blocks

As a packet from a host on the internal network reaches the NAT gateway, its source address and source TCP/UDP port number are recorded. The address and port number is translated to the public IP address and an unused port number on the public interface. When the Internet host replies to the internal host's packet, it is addressed to the NAT gateway's external IP at the translation port number. The NAT gateway searches its tables and makes the opposite changes it made to the outgoing packet. NAT then forwards the reply packet to the internal host.

Translation of ICMP packets happens in a similar fashion, but without the source port modification.

NAT can be used in static and dynamic modes. Static NAT masks the private IP addresses by translating each internal address to a unique external address. Dynamic NAT translates all internal addresses to one or more external addresses.

26.1.4. Port Forwarding

Port forwarding, also known as redirection, allows traffic coming from the Internet to be sent to a host behind the NAT gateway.

Previous examples have described the NAT process when connections are made from the intranet to the Internet. In those examples, addresses and ports were unambiguous.

When connections are attempted from the Internet to the intranet, the NAT gateway will have multiple hosts on the intranet that could accept the connection. It needs additional information to identify the specific host to accept the connection.

Suppose that two hosts, 192.168.1.10 and 192.168.1.20 are located behind a NAT gateway having a public interface of 213.18.101.62. When a connection request for http port 80 arrives at 213.18.101.62, the NAT gateway could forward the request to either of the hosts (or could accept it itself). Port forwarding configuration could be used to redirect the requests to port 80 to the first host.

Port forwarding can also remap port numbers. The second host may also need to answer http requests. As connections to port 80 are directed to the first host, another port number (such as 8080) can be dedicated to the second host. As requests arrive at the gateway for port 8080, the gateway remaps the port number to 80 and forwards the request to the second host.

Finally, port forwarding can take the source address into account. Another way to solve the above problem could be to dedicate two hosts 200.0.0.1 and 200.0.0.2 and have the NAT gateway forward requests on port 80 from 200.0.0.1 to 192.168.1.10 and from 200.0.0.2 to 192.168.1.20.

26.2. Firewall Quick Setup

For users familiar with the firewall the following will serves as a reminder of how to build the firewall. New users may wish to read Section 26.3, "Firewall Terminology And Concepts" before continuing.

1. Logically partition your network into zones. Will you establish a DMZ? Will all Ethernet interfaces need to forward traffic to the public network? Which interfaces are to be treated in a similar fashion?
2. Assign your interfaces to the zones. If using T1/E1, have you created your T1/E1 interfaces prior to building the firewall?
3. Set the default policies for traffic from zone to zone to be as restrictive as possible. Has the local zone been blocked from connecting to the DMZ or firewall? Does the DMZ or firewall need to accept connections? Which connections should be dropped and which reset? What logs are kept?
4. How is the network interface IP assigned, i.e. dynamically or statically? Do hosts at the central site need to know the local address?

5. If your network interface IP is dynamically assigned, configure masquerading.
6. If your network interface IP is statically assigned, configure Source Network address Translation (SNAT). If a sufficient number of IP addresses are provided by the ISP, static NAT can be employed instead.
7. If your hosts must accept sessions from the Internet, configure the rules file to support Destination Network address Translation (DNAT). Which hosts need to accept connections, from whom and on which ports?
8. Configure the rules file to override the default policies. Have external connections been limited to approved IP address ranges. Have all but the required protocols been blocked?
9. If you are supporting a VPN, add additional rules.
10. Validate the configuration using the method outlined in Section 26.5.2, “Working with Firewall Configurations”.
11. Activate the firewall. It is recommended to run a port scan of the firewall after activation and verify that any defined logging is functioning as expected.

26.3. Firewall Terminology And Concepts

This section provides background on various firewall terms and concepts. References are made to the section where configuration applies.

26.3.1. Zones

A network zone is a collection of interfaces, for which forwarding decisions are made, for example:

Name	Description
net	The Internet
loc	Your Local Network
dmz	Demilitarized Zone
fw	The firewall itself
vpn1	IPSec connections on w1ppp
vpn2	IPSec connections on w2ppp

Table 26.2. Network Zones

New zones may be defined at any time. For example, if all of your Ethernet interfaces are part of the local network zone, disallowing traffic from the Internet zone to the local zone will disallow it to all Ethernet interfaces. If you wanted some interfaces (but not others) to access the Internet, you could create another zone.

26.3.2. Interfaces

ROX™ Firewall interfaces are simply the LAN and WAN interfaces available to the router. You must place each interface into a network zone.

If an interface supports more than one subnet, place the interface in zone ‘Any’ and use the zone hosts setup (see below) to define a zone for each subnet on the interface.

An example follows:

Interface	Zone
fe-3-1	loc
fe-3-2	loc
fe-3-3	Any
fe-3-4	dmz

Interface	Zone
w1ppp	net

Table 26.3. Interfaces

26.3.3. Hosts

ROX™ firewall hosts are used to assign zones to individual hosts or subnets, on an interface which handles multiple subnets. This allows the firewall to manage traffic being forwarded back out the interface it arrived on, but destined for another subnet. This is often useful for VPN setups to handle the VPN traffic separately from the other traffic on the interface which carries the VPN traffic. An example follows:

Zone	Interface	IP Address or Network
local	fe-3-3	10.0.0.0/8
guests	fe-3-3	192.168.0.0/24

Table 26.4. Hosts

26.3.4. Policy

Firewall policies are the default actions for connection establishment between different firewall zones. Each policy is of the form:

Source-zone Destination-zone Default-action

You can define a policy from each zone to each other. You may also use a wildcard zone of “all” to represent all zones.

The default action describes how to handle the connection request. There are six types of actions: ACCEPT, DROP, REJECT, QUEUE, CONTINUE and NONE. The first three are the most widely used and are described here.

When the *ACCEPT* policy is used, a connection is allowed. When the *DROP* policy is used, a request is simply ignored. No notification is made to the requesting client. When the *REJECT* policy is used, a request is rejected with an TCP RST or an ICMP destination-unreachable packet being returned to the client.

An example should illustrate the use of policies.

Source Zone	Destination Zone	Policy
loc	net	ACCEPT
net	all	DROP
all	all	REJECT

Table 26.5. Policies

The above policies will:

- Allow connection requests *only* from your local network to the Internet. If you want to allow requests from a ROX™ console to the Internet, add a policy of ACCEPT fw zone to the net zone.
- Drop (ignore) all connection requests from the Internet to your firewall or local network, and
- Reject all other connection requests.

Note that a client on the Internet probing the TCP/UDP ports will receive no responses and will not be able to detect the presence of the router. A host in the local network will fail to connect to the router, but will receive a notification.

Note that the order of policies is important. If the last rule of this example were entered first then no connections at all would be allowed.

26.3.5. Masquerading and SNAT

Masquerading and Source NAT (SNAT) are forms of dynamic NAT.

Masquerading substitutes a single IP address for an entire internal network. Use masquerading when your ISP assigns you an IP address dynamically at connection time.

SNAT substitutes a single address or range of addresses that have been assigned by your ISP. Use SNAT when your ISP assigns you one or more static IP addresses that you wish to use for one or more internal hosts.

Interface Subnet Address Protocol Port(s)

Interface is the outgoing (WAN or Ethernet) interface and is usually your Internet interface.

Subnet is the subnet that you wish to hide. It can be an interface name (such as fe-3-1) or a subnetted IP address.

Address is an (optional IP) address that you wish to masquerade as.

The presence of the Address field determines whether masquerading or SNAT is being used. Masquerading is used when only Interface and Subnet are present. SNAT is used when Interface, Subnet and Address are present.

Protocol (optionally) takes on the name of protocols (e.g. tcp, udp) that you wish to masquerade.

Ports (optionally) takes on the ports to masquerade when protocol is set to tcp or udp. These can be raw port numbers or names as found in file /etc/services.

Some examples should illustrate the use of masquerading:

Rule	Interface	Subnet	Address	Protocol	Ports
1	fe-3-1	fe-3-2			
2	ppp+	fe-3-2	66.11.180.161		
3	ppp+	192.168.0.0/24	66.11.180.161		
4	w1ppp	fe-3-1	100.1.101.16		
5	w1ppp	fe-3-1	100.1.101.16	tcp	smtp

Table 26.6. Masquerading Examples

1. In this masquerading rule, port fe-3-2 is connected to the local network and fe-3-1 is connected to a DSL modem. Traffic from the subnet handled by fe-3-2 should be translated to whatever IP is assigned to the modem. Internet clients will not be able to determine the router's public address unless some form of dynamic DNS is employed.
2. In this SNAT rule, a static address of 66.11.180.161 is acquired from the ISP. Traffic from the subnet handled by fe-3-2 should be translated to 66.11.180.161 as it sent to the Internet over ppp. The + at the end of "ppp+" causes the ROX™ firewall to match any ppp interface.
3. This example is much the same as the previous one only the subnet is explicitly described, and could include traffic from any of the Ethernet ports.
4. In this SNAT rule, traffic from the subnet handled by only port fe-3-1 should be translated to 100.1.101.16 as it sent to the Internet on t1/e1 port w1ppp.
5. This example is much the same as the previous one excepting that only smtp from fe-3-1 will be allowed.

26.3.6. Rules

The default policies can completely configure traffic based upon zones. But the default policies cannot take into account criteria such as the type of protocol, IP source/destination addresses and the need to perform special actions such as port forwarding. The firewall rules can accomplish this.

The ROX™ firewall rules provide exceptions to the default policies. In actuality, when a connection request arrives, the rules file is inspected first. If no match is found then the default policy is applied. Rules are of the form:

Action Source-Zone Destination-Zone Protocol Destination-Port Source-Port Original-Destination-IP Rate-Limit User-Group

Actions are ACCEPT, DROP, REJECT, DNAT, DNAT-, REDIRECT, REDIRECT-, CONTINUE, LOG and QUEUE. The DNAT-, REDIRECT-, CONTINUE, LOG and QUEUE actions are not widely used used and are not described here.

Action	Description
ACCEPT	Allow the connection request to proceed.
DROP	The connection request is simply ignored. No notification is made to the requesting client.
REJECT	The connection request is rejected with an RST (TCP) or an ICMP destination-unreachable packet being returned to the client.
DNAT	Forward the request to another system (and optionally another port).
REDIRECT	Redirect the request to a local tcp port number on the local firewall. This is most often used to “remap” port numbers for services on the firewall itself.

Table 26.7.

The remaining fields of a rule are as described below:

Rule Field	Description
Action	The action as described in the previous table.
Source-Zone	The zone the connection originated from.
Destination-Zone	The zone the connection is destined for.
Protocol	The tcp or udp protocol type.
Destination-Port	The tcp/udp port the connection is destined for.
Source-Port	The tcp/udp port the connection originated from.
Original-Destination-IP	The destination IP address in the connection request as it was received by the firewall.
Rate-Limit	A specification which allows the rate at which connections are made to be limited.

Table 26.8.

Some examples will illustrate the power of the rules file:

Rule	Action	Source-Zone	Destination-Zone	Protocol	Dest-Port	Source-Port	Original-Destination-IP
1	ACCEPT	net:204.18.45.0/24	fw				
2	DNAT	net	loc:192.168.1.3	tcp	ssh, http		
3	DNAT	net:204.18.45.0/24	loc:192.168.1.3	tcp	http	-	130.252.100.69
4	ACCEPT	fw	net	icmp			
5	ACCEPT	net:204.18.45.0/24	fw	icmp	8		

Table 26.9.

1. This rule accepts traffic to the firewall itself from the 204.18.45.0/24 subnet. If the default policy is to drop all requests from net to the firewall, this rule will only accept traffic from the authorized subnet.
2. This rule forwards all ssh and http connection requests from the Internet to local system 192.168.1.3.

3. This rule forwards http traffic from 204.18.45.0/24 (which was originally directed to the firewall at 130.252.100.69) to the host at 192.168.1.3 in the local zone. If the firewall supports another public IP address (e.g. 130.252.100.70), a similar rule could map requests to another host.
4. and 5. These rules allow the firewall to issue icmp requests to the Internet and to respond to icmp echo requests from the authorized subnet.

Each of the Source and Destination zones may use one of the defined zone names, or one may select "Other..." and specify a zone name in the text field to the right. Both Source and Destination may be further qualified using the Only hosts in zone with addresses fields. Multiple comma-separated subnet, IP, or MAC addresses may be specified in the following way:

- IP subnet: 192.168.1.0/24
- IP address: 192.168.1.1
- IP address range: 192.168.1.1-192.168.1.25
- MAC address: ~00-A0-C9-15-39-78

26.4. Configuring The Firewall And VPN

26.4.1. Policy-based Virtual Private Networking

Begin configuration by creating local, network and vpn zones. Identify the network interface that carries the encrypted IPSec traffic and make this interface part of zone “ANY” in the interfaces menu as it will be carrying both traffic for both zones.

Visit the Host menu and, for the network interface that carries the encrypted IPSec traffic, create a zone host with zone VPN, the correct subnet and the IPSec zone option checked. If you plan to have VPN tunnels to multiple remote sites ensure that a zone host entry exists for each (or collapse them into a single subnet). Create another zone host for the same interface with a network zone, using a wider subnet mask such as 0.0.0.0/0. It is important that the vpn zone be declared before the net zone since the more specific vpn zone subnet must be inspected first.

Host Zone	Interface	Subnet	IPSec Zone?
vpn	w1ppp	192.168.1.0/24	Yes
net	w1ppp	0.0.0.0/0	No

Table 26.10.

The IPSec protocol operates on UDP port 500 and using protocols Authentication Header (AH) and Encapsulating Security Payload (ESP) protocols. The firewall must accept this traffic in order to allow IPSec.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	net	fw	ah	
ACCEPT	net	fw	esp	
ACCEPT	net	fw	udp	500

Table 26.11.

IPSec traffic arriving at the firewall is directed to openswan, the IPSec daemon. Openswan then decrypts the traffic and forwards it back to the ROX™ firewall on the same interface that originally received it. You will also need a rule to allow traffic to enter from this interface.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	vpn	loc		

Table 26.12.

26.4.2. Virtual Private Networking to a DMZ

If the firewall is to pass the VPN traffic through to another device (e.g. a VPN device in a DMZ) then establish a DMZ zone and install the following rules.

Action	Source-Zone	Destination-Zone	Protocol	Dest-Port
ACCEPT	net	dmz	ah	
ACCEPT	net	dmz	esp	
ACCEPT	net	dmz	udp	500
ACCEPT	dmz	net	ah	
ACCEPT	dmz	net	esp	
ACCEPT	dmz	net	udp	500

Table 26.13.

26.5. Firewall Configuration

All firewall fields accept only alphanumeric characters, excluding spaces. Do not use punctuation or other special characters in these fields.

Figure 26.1. Security Menu

The Security menu is a top-level menu that is accessible from the main menu. Items used to configure network security can be accessed from this menu.

Figure 26.2. Firewall Description table

Name

Synopsis: *string*

Description

Synopsis: *string*

An optional description string

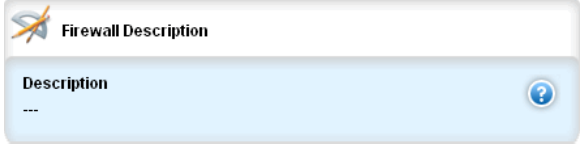

Figure 26.3. Firewall Description form

To display the Firewall form, navigate to *security/firewall* and then click on the submenu representing the configured firewall (for example, *firewall1*).

26.5.1. Adding a Firewall

To add a firewall, enter edit private mode, navigate to */security/firewall/fwconfig*, and click <Add fwconfig>.

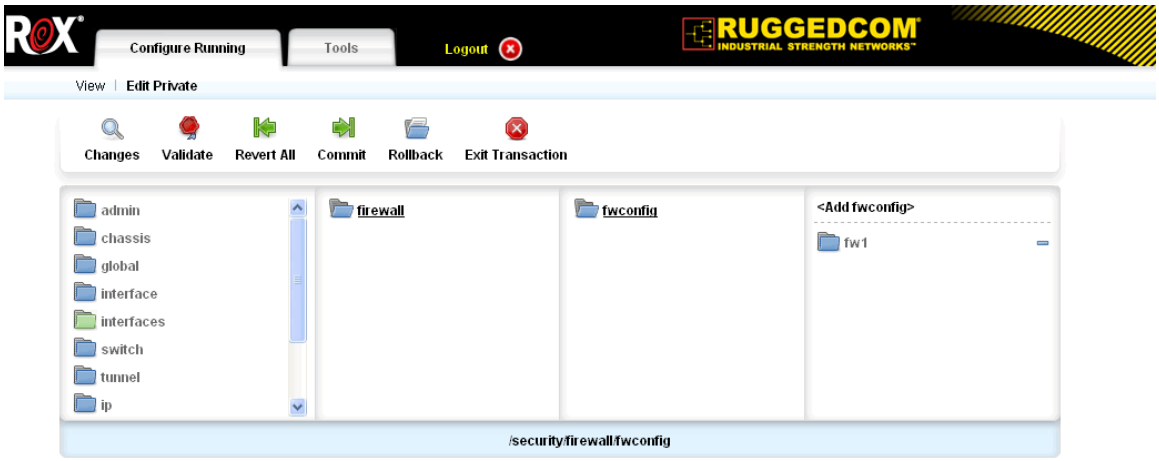

Figure 26.4. Adding a Firewall

In the *Key settings* form, enter a name for the firewall and click *Add* field.

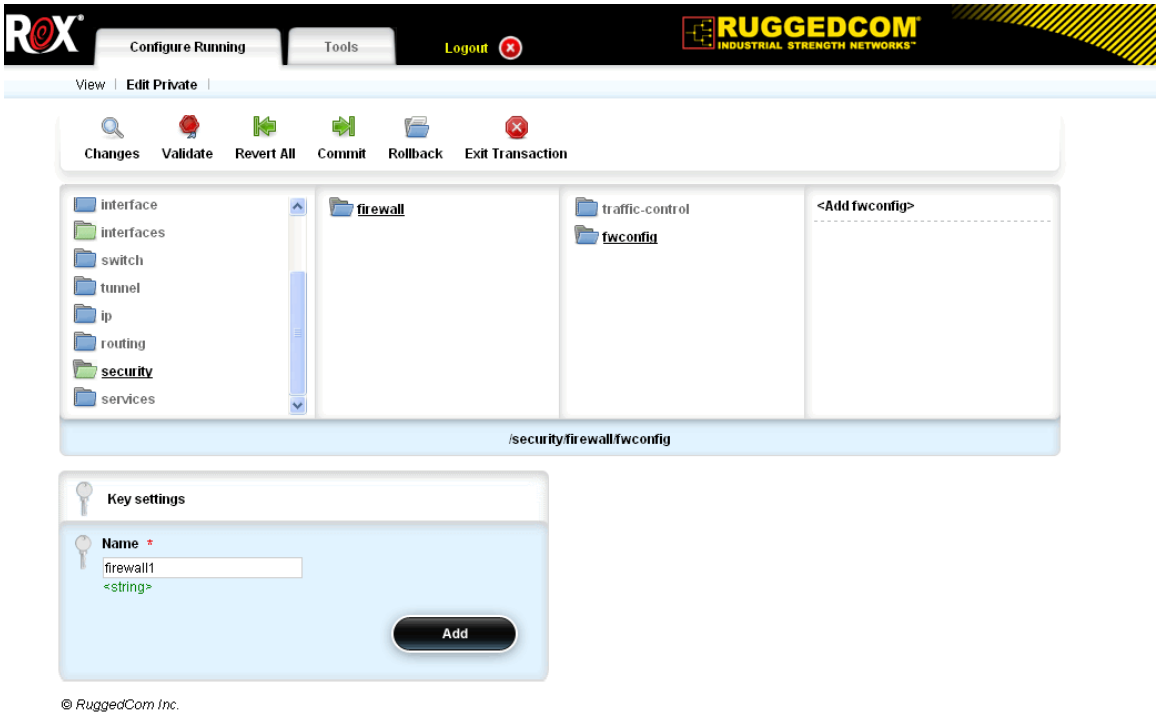

Figure 26.5. Naming a Firewall

After adding the firewall, the fwmasq, fwnat, fwrule, fwpolicy, fwinterface, fwhost, and fwzone submenus appear. To configure these areas, click on a submenu.

Figure 26.6. Firewall Submenus

26.5.2. Working with Firewall Configurations

The ROX™ firewall configuration system allows a network security administrator to work on one or more inactive firewall configurations while another is active and installed on the system. Section 26.5.2.1, "Typical Use Case" illustrates how to use the ROX™ firewall configuration system.

Control of the firewall configuration is achieved by using the three variables in the *Firewall Configuration* form, below:

Figure 26.7. Firewall Configuration form

Enable active configuration

Enables/disables the firewall configuration specified in active-config

Specify work configuration

Synopsis: *string*

The current work firewall is specified here.

Specify active configuration

Synopsis: *string*

The current active firewall is specified here

26.5.2.1. Typical Use Case

The following set of steps illustrates the configuration and maintenance of a set of firewall rules on an active ROX™ firewall system:

1. On an unconfigured system, begin configuring a set of firewall rules by giving the firewall a name: 'fw1', adding zones, interfaces, etc. At each commit at this stage, configuration data is saved but no validation is performed.
2. In order to validate the 'fw1' firewall configuration in progress, set the *work-config* variable to the name: 'fw1' and commit the changes. The system validates the firewall configuration named 'fw1' and displays the results. Note that the configuration in progress is saved whether or not the

validation succeeds. A configuration in progress may be validated in this way at any time without affecting an active firewall configuration.

3. After ‘fw1’ has been verified, it may be made active in the system by setting the *active-config* variable to the name: ‘fw1’, setting *firewall-enable* and committing the changes. The system validates the active firewall configuration and if it finds no errors, it activates the ‘fw1’ firewall configuration.
4. While the ‘fw1’ firewall configuration is active, you might wish to make changes without altering the live configuration. Using the CLI, copy the firewall configuration named ‘fw1’ to ‘fw2’. Change the *work-config* variable to ‘fw2’. Any configuration changes made to ‘fw2’ are validated when you commit your changes, and any errors in ‘fw2’ are displayed. An alternate configuration may be modified and validated in this way at any time without affecting an active firewall configuration.
5. Alternatively, while the ‘fw1’ firewall configuration is active, you might wish to make changes to the live configuration. Any changes made to a configuration that is defined as ‘active-config’ and ‘enable’ will be reflected on the live configuration currently running on the system, pending a successful validation. For instance, work-config can be a configuration named ‘fw2’ while active-config is ‘fw1’ and enabled. Modifying fw1 in this case will, upon successful validation, update the running configuration to reflect the changes.

26.5.3. Zone Configuration

Zones are collections of interfaces, for which forwarding decisions are made. They are made of different networks reachable from this system, defined by name and type of zone.

Figure 26.8. Zone table

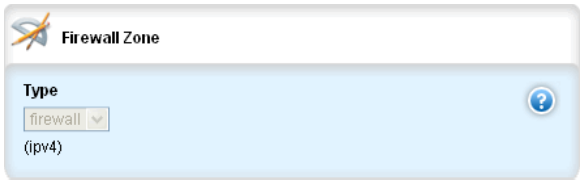

Figure 26.9. Zone form

This form allows you to add, delete and configure zones. New zones can be added by entering the Edit Private mode and then adding zones.

Name

Synopsis: *A string*

A unique name to assign to this zone. Be sure to *also create* a zone called **fw** of type *firewall*.

Type

Synopsis: *string - one of the following keywords { firewall, ipsec, ipv4 }*

Default: *ipv4*

Zone types are: firewall, ipv4 or ipsec

description

Synopsis: *string*

(Optional) The description string for this zone

26.5.4. Interface Configuration

Figure 26.10. Main Interface Settings table

interface

Synopsis: *string*

Currently active or not - add '+' for same interfaces: ppp+

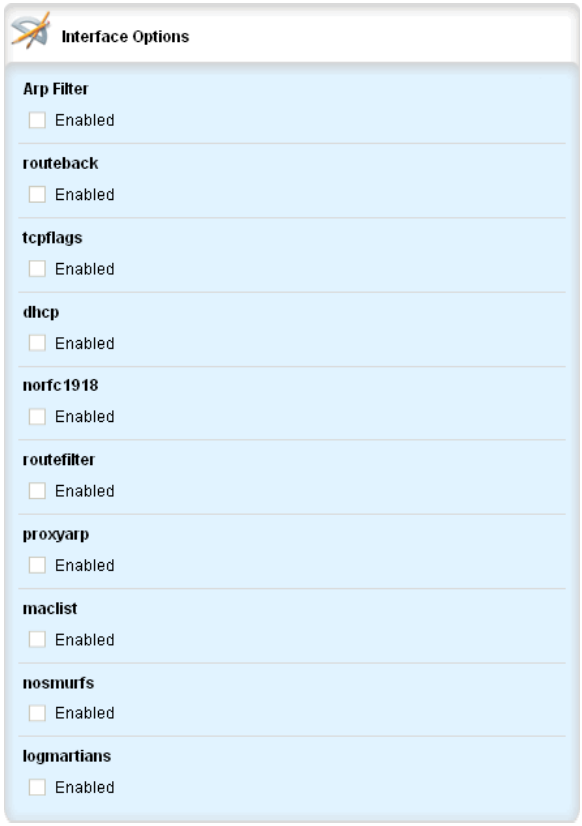

Figure 26.11. Interface Options form

Arp Filter

Responds only to ARP requests for configured IP addresses

routeback

Allow traffic on this interface to be routed back out that same interface

tcpflags

Illegal combinations of TCP flags dropped and logged at info level

dhcp

Allows DHCP datagrams to enter and leave the interface

norfc1918

Not currently implemented

routefilter

Enables route filtering

proxyarp

Enables proxy ARP

maclist

Not currently implemented

nosmurfs

Packets with broadcast address as source are dropped and logged at info level

logmartians

Enables logging of packets with impossible source addresses

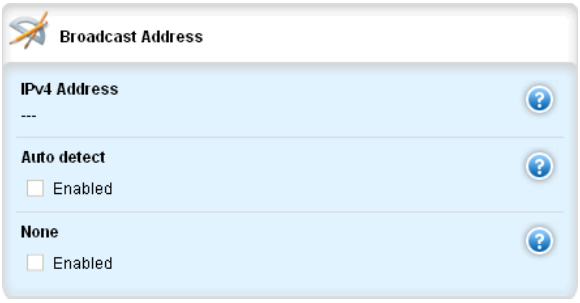

Figure 26.12. Broadcast Address form

broadcast-addr

(Optional) A broadcast address

26.5.5. Host Configuration

Hosts are used to assign zones to individual hosts or subnets, on an interface which handles multiple subnets.

Figure 26.13. Main Host Settings table

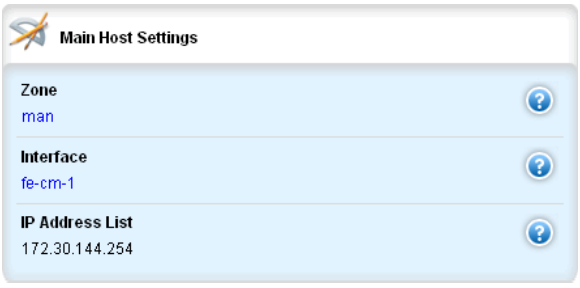

Figure 26.14. Main Host Settings form

name

Synopsis: *string*

The name of a host configuration entry

Zone

Synopsis: *A string*

A pre-defined zone

Interface

Synopsis: *A string*

A pre-defined interface to which optional IPs and/or networks can be added

IP Address List

Synopsis: *string*

(Optional) Additional IP addresses or networks - comma separated

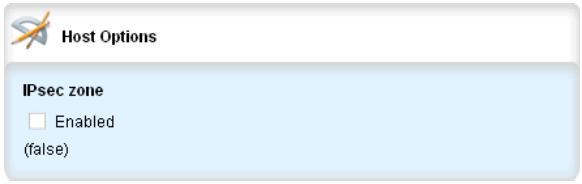

Figure 26.15. Host Options form

IPsec zone

Synopsis: *boolean*

Default: *false*

26.5.6. Policies

Figure 26.16. Main Policy Settings table

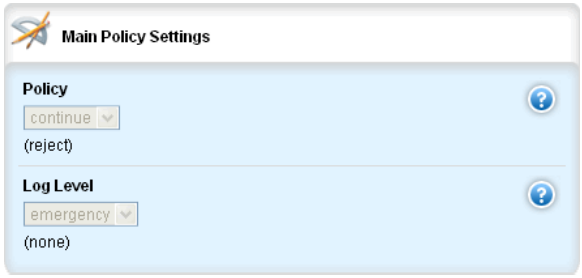

Figure 26.17. Main Policy Settings form

Default actions for connection establishment between different zones. Configuration of the default actions for traffic between different firewall zones. They can be overridden for particular hosts or types of traffic by defining specific rules.

Policy Name

Synopsis: *string*

Enter a name tag for this policy

Policy

Synopsis: *string - one of the following keywords { continue, reject, drop, accept }*

Default: *reject*

A default action for connection establishment between different zones.

Log Level

Synopsis: *string - one of the following keywords { emergency, alert, critical, error, warning, notice, info, debug, none }*

Default: *none*

(Optional) Whether or not logging will take place and at which logging level.

description

Synopsis: *string*

(Optional) The description string for this policy

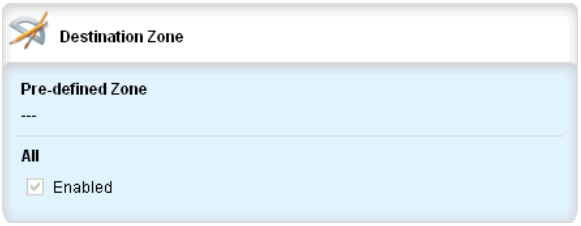

Figure 26.18. Destination Zone form

destination-zone

The zone in which the request terminates. Enter a destination zone configuration by specifiying a zone. Please choose either a pre-defined zone or all.

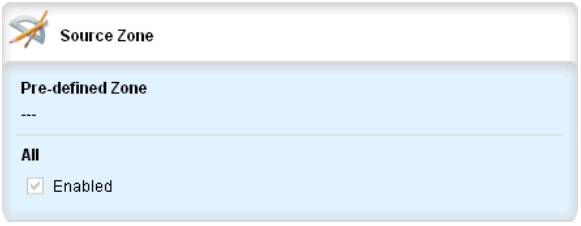

Figure 26.19. Source Zone form

source-zone

The zone from which the request originates. Enter a source zone configuration by specifiying a zone. Please choose either a pre-defined zone or all.

26.5.7. Network Address Translation

Configuration of one-to-one Network Address Translation (NAT). The static network address translation entries in this table can be used to set up a one-to-one correspondence between an external address on your firewall and an RFC1918 address of a host behind the firewall. Static NAT is often used to allow connections to an internal server from outside the network.

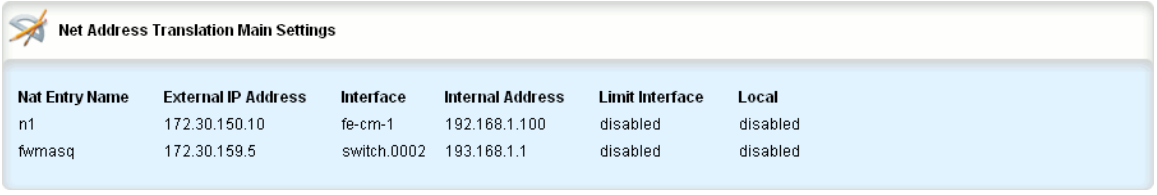

Figure 26.20. Net Address Translation Main Settings table

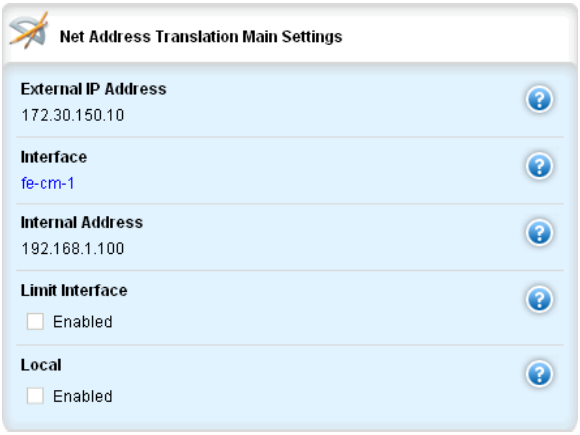

Figure 26.21. Net Address Translation Main Settings form

Nat Entry Name

Synopsis: *string*

Enter a name for this NAT entry

External IP Address

Synopsis: *IPv4 address in dotted-decimal notation*

The external IP Address (must not be a DNS name)

Interface

Synopsis: *A string*

Interfaces that have the EXTERNAL address

Internal Address

Synopsis: *IPv4 address in dotted-decimal notation*

The internal address (must not be a DNS Name)

Limit Interface

Translation only effective from the defined interface (default: all interfaces)

Local

Translation effective from the firewall system (default: disabled)

description

Synopsis: *string*

(Optional) The description string for this nat entry

26.5.8. IP Masquerading

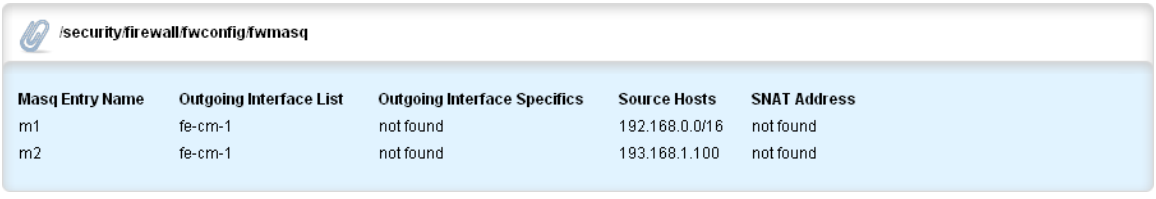

Figure 26.22. FWMasq table

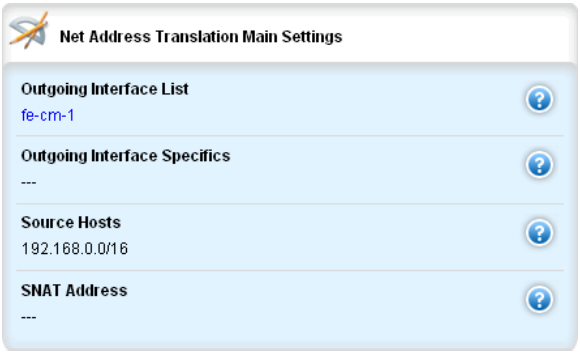

Figure 26.23. Net Address Translation Main Settings form

Masquerading substitutes a single IP address for an entire internal network

Masq Entry Name

Synopsis: *string*

A name for this masquerading configuration entry

Outgoing Interface List

Synopsis: *string*

An outgoing interfacelist - usually the internet interface

Outgoing Interface Specifics

Synopsis: *string*

(Optional) An outgoing interface list - specific destinations IP for the out-interface

Source Hosts

Synopsis: *string*

Subnet range or comma-separated list of hosts (IPs)

SNAT Address

Synopsis: *string*

(Optional) By specifying an address here, SNAT will be used and this will be the source address

description

Synopsis: *string*

(Optional) The description string for this masq entry

26.5.9. Rules

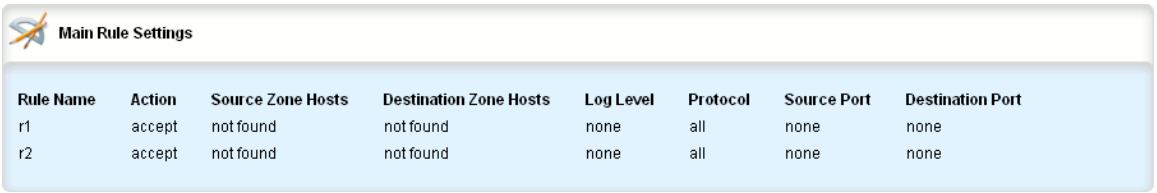
Main Rule Settings
Rule Name Action Source Zone Hosts Destination Zone Hosts Log Level Protocol Source Port Destination Port
r1 accept not found not found none all none none
r2 accept not found not found none all none none

Figure 26.24. Main Rule Settings table

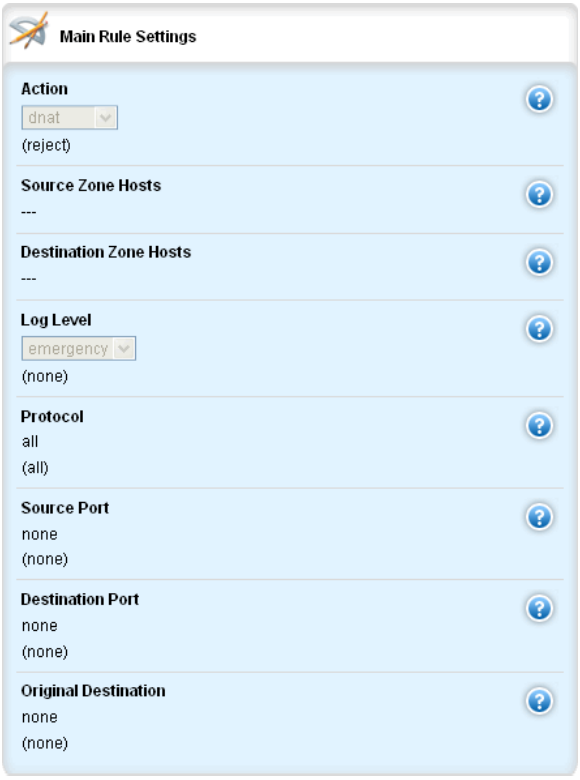

Figure 26.25. Main Rule Settings form

Rules are to establish exceptions to the default policies. This table lists exceptions to the default policies for certain types of traffic, sources or destinations. The chosen action will be applied to packets matching the chosen criteria instead of the default.

Rule Name

Synopsis: *string*

Enter a unique name that identifies this rule.

Action

Synopsis: *string - one of the following keywords { dnat, dnat-, redirect, continue, reject, drop, accept }*

Default: *reject*

The final action to take on incoming packets matching this rule.

Destination Zone Hosts

Synopsis: *string*

(Optional) Add comma-separated host IPs to the destination-zone - may include :port for DNAT or REDIRECT

Log Level

Synopsis: *string - one of the following keywords { emergency, alert, critical, error, warning, notice, info, debug, none }*

Default: *none*

(Optional) Whether or not logging will take place and at which logging level.

Protocol

Synopsis: *string*

Synopsis: *string - one of the following keywords { all, icmp, udp, tcp }*

Default: *all*

The protocol to match for this rule.

Source Port

Synopsis: *string*

Synopsis: *string - one of the following keywords { none, Related, Any }*

Default: *none*

(Optional) The tcp/udp port the connection originated from.

Destination Port

Synopsis: *string*

Synopsis: *string - one of the following keywords { none, Related, Any }*

Default: *none*

(Optional) The tcp/udp port the connection is destined for.

Original Destination

Synopsis: *string*

Synopsis: *string - the keyword { None }*

Default: *none*

(Optional) The destination IP address in the connection request as it was received by the firewall.

description

Synopsis: *string*

(Optional) The description string for this rule

Figure 26.26. Source Zone form

Source Zone Hosts

Synopsis: *string*

(Optional) Add comma-separated host IPs to a predefined source-zone

Figure 26.27. Destination Zone form

destination-zone

synopsis: string

(Optional) Add comma-separated host IPs to the destination-zone - may include :port for DNAT or REDIRECT

27. Traffic Control

Traffic Control (TC) is a firewall subsystem managing the amount of bandwidth per network interface that different types of traffic are permitted to use. For a traffic control configuration to work, a firewall must be configured.

A ROX™ system allows up to 4 different firewall configurations, enabling you to quickly change between configurations. You can quickly assess different configurations without needing to save and reload any part of the configuration. In contrast, there is only one traffic control configuration. When enabled, a traffic control configuration is used with the current firewall configuration. A current firewall configuration is defined as one that is specified in either work-config and/or active-config. It does not have to be enabled to be validated.

A TC configuration can be seen as an additional configuration that goes along with a current firewall configuration. To actually add a TC configuration, it has to be enabled by setting the /qos/traffic-control/ Basic or Advanced Configuration Modes variable. Only at that point will a TC configuration be added to the current firewall configuration.

27.1. Traffic Control Modes

Traffic Control functions are divided into two modes: *basic mode* and *advanced mode*. The basic mode contains functions with basic traffic control configuration parameters. The advanced mode contains functions with advanced traffic control configuration parameters. The two modes cannot be accessed simultaneously. Only the mode that is currently configured can be accessed.

27.1.1. Traffic Control Basic (basic-configuration) Configuration Mode

Basic-configuration mode offers a limited set of options and parameters. Use this mode to set the outgoing bandwidth for an interface, the interface priority (high, medium, low), and some simple traffic control characteristics. Basic traffic shaping affects traffic identified by protocol, port number, address, and interface. Note that some of these options are mutually exclusive; refer to the information given for each option.

In basic-configuration mode, a packet is categorized based on the contents of its TOS field if it does not match any of the defined bands.

27.1.2. Traffic Control Advanced (advanced-configuration) Configuration Mode

In advanced-configuration mode, each interface to be managed is assigned a total bandwidth that it should allow for incoming and outgoing traffic. Classes are then defined for each interface, each with its own minimum assured bandwidth and a maximum permitted bandwidth. The combined minimum of the classes on an interface must be no more than the total outbound bandwidth specified for the interface. Each class is also assigned a priority, and any bandwidth left over after each class has received its minimum allocation (if needed) will be allocated to the lowest priority class up until it reaches its maximum bandwidth, after which the next priority is allocated more bandwidth. When the specified total bandwidth for the interface is reached, no further packets are sent, and any further packets may be dropped if the interface queues are full.

Packets are assigned to classes on the outbound interface based on either a mark assigned to the packet, or the ToS (type of service) field in the IP header. If the ToS field matches a defined class, then the packet is allocated to that class. Otherwise, it is allocated to any class that matches the mark assigned to the packet, and if no class matches the mark, then the packet is assigned to the default class.

Marks are assigned to packets either by the TC Rules based on any of a number of parameters, such as IP address, port number, protocol, packet length, and so on.

27.1.2.1. Traffic Control Example

The goal of this example is to operate T1 port at 1.5Mbit/s and ensure that UDP source port 20000 traffic gets at least half the bandwidth, while ICMP and TCP ACK packets should have high priority, HTTP traffic gets at least 20% and at most 50%, and all other traffic should get what is left over but only up to 50% of the bandwidth.

The three TC menus would be configured as follows:

27.1.2.1.1. TC Interfaces

Interface	Inbound bandwidth	Outbound bandwidth
te1-3-1c24ppp	1500kbit	1500kbit

Table 27.1. TC Interfaces

27.1.2.1.2. TC Classes

Interface	Mark	Minimum	Maximum	Priority	Options
te1-3-1c24ppp	1	full/2	full	0	
te1-3-1c24ppp	2	28bps	full	1	tcp-ack
te1-3-1c24ppp	3	full/5	full*5/10	2	
te1-3-1c24ppp	4	28bps	full*5/10	3	default

Table 27.2. TC Classes

27.1.2.1.3. TC Rules

Mark	Source	Destination	Protocol	Source Port	Dest Port	Test	Length	TOS
2	Any	Any	ICMP	Any	Any	Any	Any	Any
RESTORE	Any	Any	Any	Any	Any	0	Any	Any
CONTINUE	Any	Any	Any	Any	Any	!0	Any	Any
1	Any	Any	UDP	20000	Any	Any	Any	Any
3	Any	Any	TCP	Any	80	Any	Any	Any
4	Any	Any	Any	Any	Any	0	Any	Any
SAVE	Any	Any	Any	Any	Any	!0	Any	Any

Table 27.3. TC Rules

The rules first check non connection-based protocol rules (ICMP in this case) in order to assign a mark. For any packet that is still not marked, we attempt to restore a saved mark for the connection. If at this point the packet has a mark set, we stop checking rules (CONTINUE) since it is either ICMP or a packet from an existing connection which we have already assigned a mark. If still no mark is assigned, it must be a new connection so we process the packet through all the remaining rules to determine the mark it should receive. At the end, we save the new mark to the connection so that any further packets for the connection do not have to go through all the rules again, in order to save processing resources. We mark all packets with no other matching rule to 4 since that represents the default class (as defined in TC Classes). This allows explicit traffic control of even unspecified network connections.

27.2. Traffic Control Configuration

Figure 27.1. Traffic-Control menu

To display the Traffic Control menu, navigate to *qos/traffic-control*.

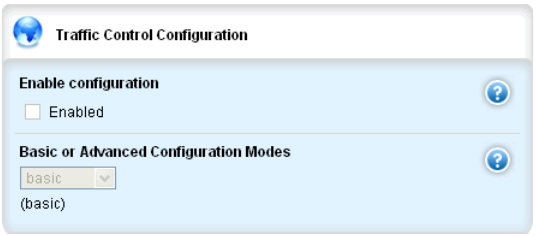

Figure 27.2. Traffic Control Configuration form

The Traffic Control Configuration form appears on the same screen as the Traffic Control menu.

Enable configuration

Enables/disables traffic control (TC) for the current firewall configuration. The current firewall configuration is the one that is committed. When an active configuration is committed to the system, then an **enabled** TC configuration will be included. When a work configuration is comitted then the **enabled** TC configuration will be included in the work config. **A TC configuration needs a firewall configuration to operate**.

Basic or Advanced Configuration Modes

Synopsis: *string - one of the following keywords { advanced, basic }*

Default: *basic*

Specifies to use either 'simple' or 'advanced' configuration modes. Click again on traffic-control after making a choice.

27.2.1. Traffic Control Modes

Traffic Control functions are divided into two modes: basic-configuration mode and advanced-configuration mode. Basic-configuration mode contains functions with basic traffic control configuration parameters. Advanced-configuration mode contains functions with advanced traffic control configuration parameters. The two modes cannot be accessed simultaneously. Only the mode that is currently configured can be accessed.

It is mandatory to configure the firewall before enabling traffic control. For information on configuring the firewall, refer to Chapter 26, Firewall.

27.2.1.1. Basic-configuration Mode

To configure basic-configuration mode, follow the procedure below.

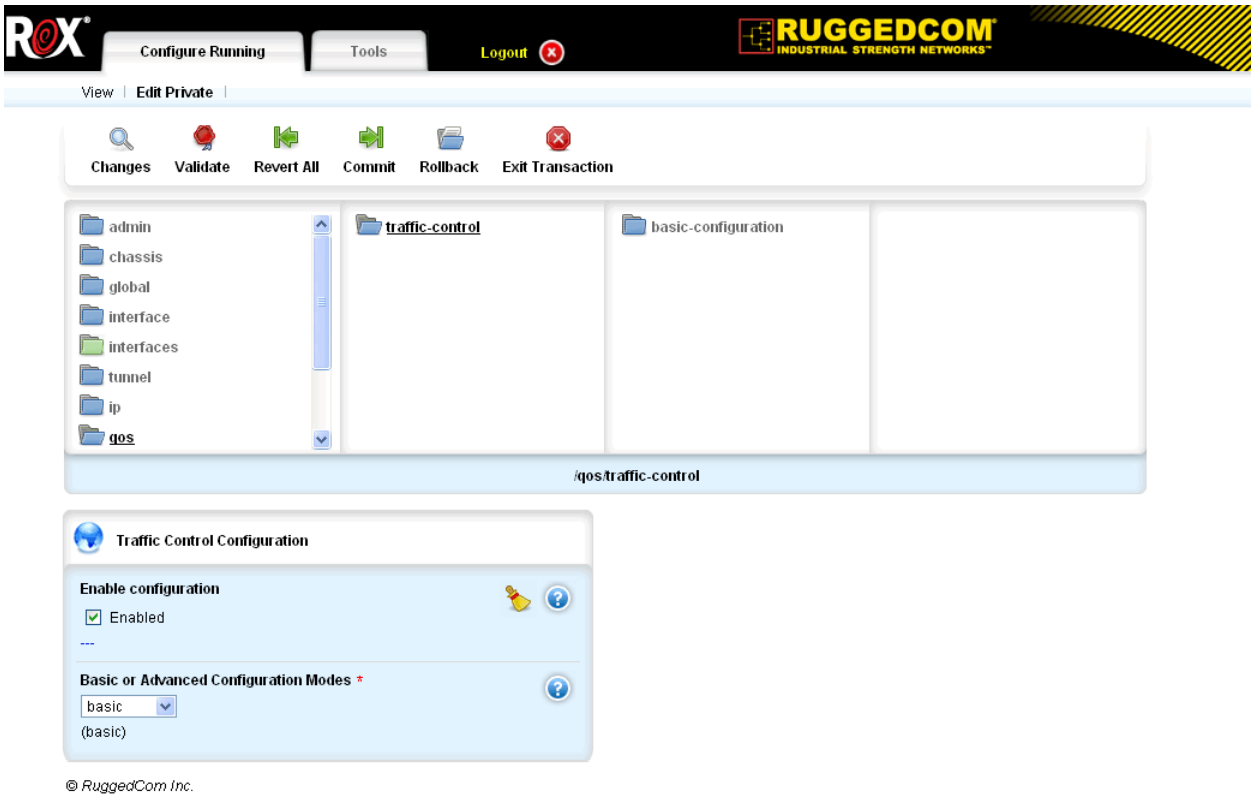

Figure 27.3. Enabling Basic-configuration Mode

Procedure 27.1. Configuring Basic-configuration Mode

1. Enter Edit Private mode.
2. Click on *qos/traffic-control*.
3. On the Traffic Control Configuration form, click **Enabled** in the *Enable configuration* field.
4. Select **basic** in the *Basic or Advanced Configuration Modes* field.
5. Click **Commit**.
6. Click **Exit Transaction**.

27.2.1.1.1. Interfaces

The interface is the network interface to which traffic shaping will apply.

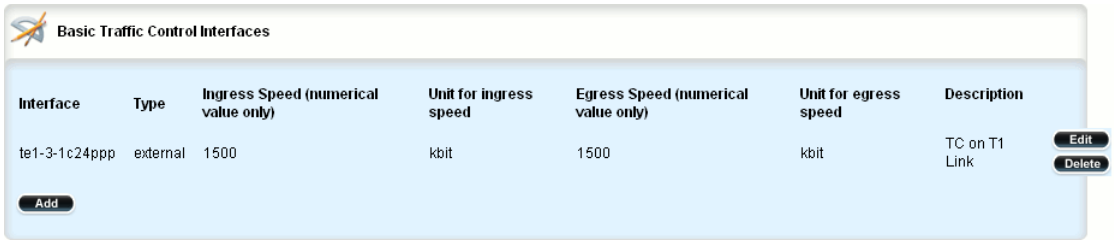

Figure 27.4. Basic Traffic Control Interfaces table

To display this table, navigate to *qos/traffic-control/basic-configuration/tcinterfaces*.

Figure 27.5. Interface to Apply Traffic Control form

To display this form, navigate to *qos/traffic-control/basic-configuration/tcinterfaces/{interface}*.

interface

Synopsis: *string*

An interface to which traffic shaping will apply

Type

Synopsis: *string - one of the following keywords { none, external, internal }*

Default: *none*

(optional) 'external' (facing toward the Internet) or 'internal'
(facing toward a local network). 'external'
causes the traffic generated by each unique
source IP address to be treated as a single
flow. 'internal' causes the traffic generated by
each unique destination IP address to be treated
as a single flow. , internal interfaces seldom
benefit from simple traffic shaping. Default: none.

Ingress Speed (numerical value only)

Synopsis: *string*

(optional) The incoming bandwidth of this interface. If incoming
traffic exceeds the given rate, received packets
are dropped randomly. When unspecified, maximum
speed is assumed. Specify only the number here.
The unit (kbps, mbps) is specified in In-unit.

Unit for ingress speed

Synopsis: *string - one of the following keywords { none, bps, mbps, mbit, kbps, kbit }*

Specifies the unit when an incoming bandwidth is specified

Egress Speed (numerical value only)

Synopsis: *unsigned short integer*

The outgoing bandwidth for this interface. Specify only the number here. The unit (kbps, mbps) is specified in Out-unit.

Unit for egress speed

Synopsis: *string - one of the following keywords { bps, mbps, mbit, kbps, kbit }*

Specifies the unit for the outgoing bandwidth

Description

Synopsis: *string*

A description for this configuration item

> *Incoming traffic is prioritized based on the matching ToS value in the IP header if nothing else is configured under a band or when IP traffic does not match with the rules specified in a band.*

> *Speed units:*
>
> - *bps = bytes per second*
> - *mbps/kbps = megabyte/kilobyte per second*
> - *mbits/kbits = megabits/kilobits per second*

> *The Egress Speed and Unit for egress speed must be configured before committing a configuration.*

27.2.1.1.2. Priorities

Priorities configure traffic priorities for basic traffic shaping.

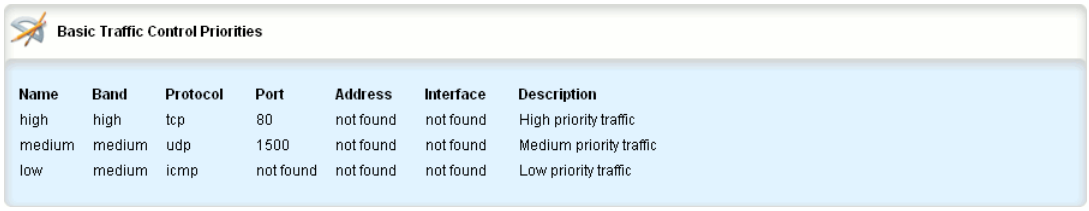
Basic Traffic Control Priorities
Name Band Protocol Port Address Interface Description
high high tcp 80 not found not found High priority traffic
medium medium udp 1500 not found not found Medium priority traffic
low medium icmp not found not found not found Low priority traffic

Figure 27.6. Basic Traffic Control Priorities table

To display this table, navigate to *qos/traffic-control/basic-configuration/tcpriorities*.

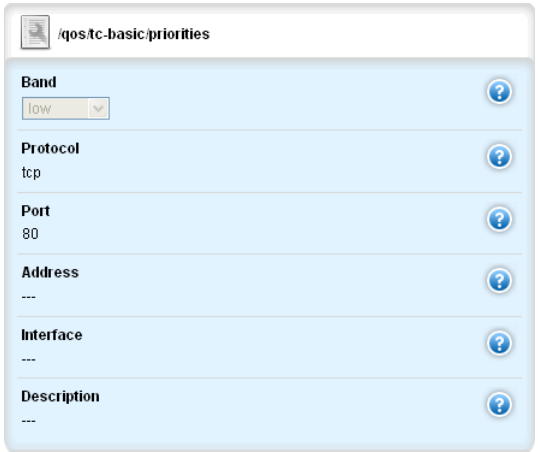

Figure 27.7. Priorities form

To display this form, navigate to *qos/traffic-control/basic-configuration/tcpriorities/{priority}*.

name

Synopsis: *string*

A distinct name for this configuration entry

band

Synopsis: *string - one of the following keywords { low, medium, high }*

Default: *medium*

Priority (band) : high, medium, low...

High band includes:

Minimize Delay (md) (0x10),

md + Minimize Monetary Cost (mmc) (0x12),

md + Maximize Reliability (mr) (0x14),

mmc+md+mr (0x16).

Medium band includes:

Normal Service (0x0),

mr (0x04),

mmc+mr (0x06),

md + Maximize Throughput (mt) (0x18),

mmc+mt+md (0x1a),

mr+mt+md (0x1c),

mmc+mr+mt+md (0x1e).

Low band includes:

mmc (0x02),

mt (0x08),

mmc+mt (0x0a),

mr+mt (0x0c),

mmc+mr+mt (0x0e).

protocol

Synopsis: *string*

Synopsis: *string - one of the following keywords { all, icmp, udp, tcp }*

(choice) Targeted protocol

port

Synopsis: *string*

(choice) Source port - can be specified **only if** protocol is tcp, udp, dccp, sctp or udplite

address

Synopsis: *string*

(choice) Source address - can be specified **only if** protocol, port and interface are not defined

interface

Synopsis: *string*

(choice) Source interface - can be specified **only if** protocol, port and address are not defined

description

Synopsis: *string*

(optional) A description for this configuration

> *For basic traffic control configurations, **Port**, **Address** and **Interface** refer to the source of the traffic.*

27.2.1.2. Advanced-configuration Mode

To configure advanced-configuration mode, follow the procedure below.

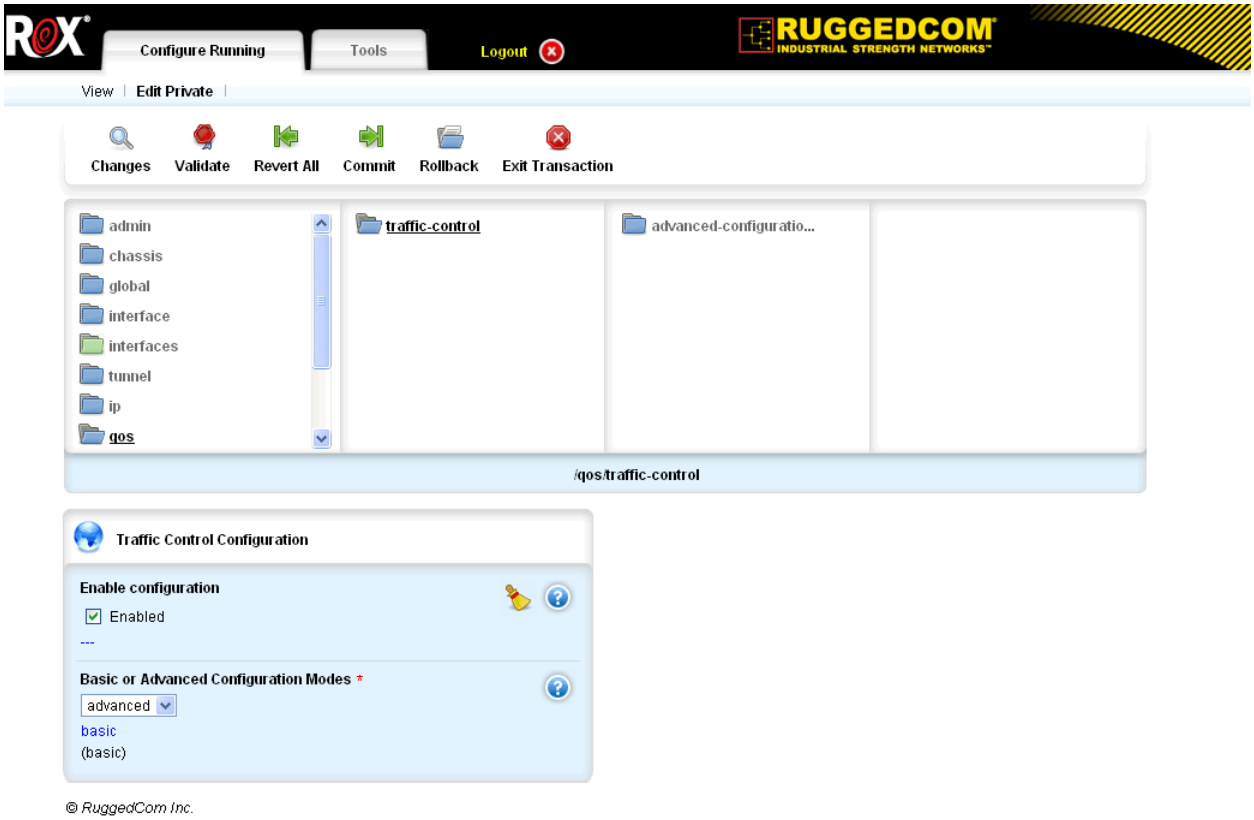

Figure 27.8. Enabling Advanced-configuration Mode

Procedure 27.2. Configuring Advanced-configuration Mode

1. Enter Edit Private mode.
2. Click on *qos/traffic-control*.
3. On the Traffic Control Configuration form, click **Enabled** in the *Enable configuration* field.
4. Select **advanced** in the *Basic or Advanced Configuration Modes* field.
5. Click **Commit**.
6. Click **Exit Transaction**.

27.2.1.2.1. Classes

Classes define classes for traffic shaping.

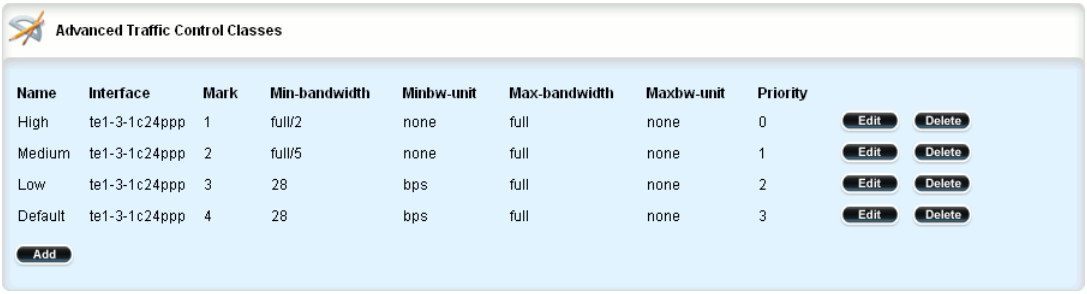
Advanced Traffic Control Classes
Name Interface Mark Min-bandwidth Minbw-unit Max-bandwidth Maxbw-unit Priority
High te1-3-1c24ppp 1 full/2 none full none 0 Edit Delete
Medium te1-3-1c24ppp 2 full/5 none full none 1 Edit Delete
Low te1-3-1c24ppp 3 28 bps full none 2 Edit Delete
Default te1-3-1c24ppp 4 28 bps full none 3 Edit Delete
Add

Figure 27.9. Advanced Traffic Control Classes table

To display this table, navigate to *qos/traffic-control/advanced-configuration/tcclasses*.

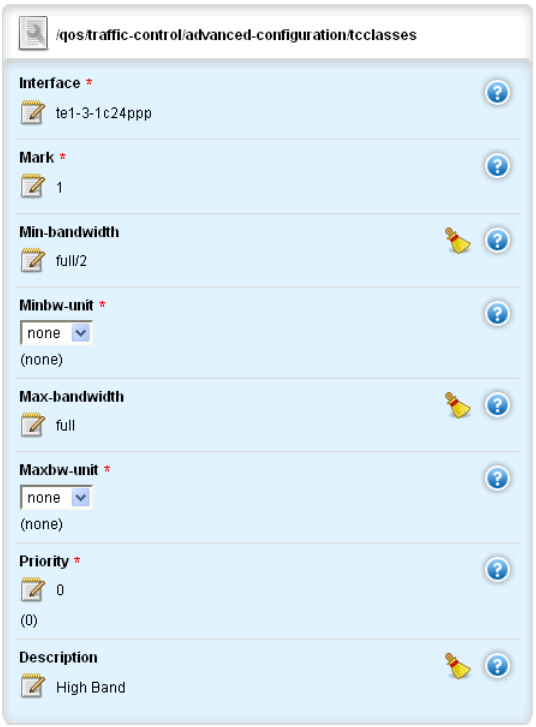

Figure 27.10. TC Classes form

To display this form, navigate to *qos/traffic-control/advanced-configuration/tcclasses/{class}*.

Note that each class is associated with exactly one network interface. Exactly one class for each interface must be designated as the default.

name

Synopsis: *string*

A name for this TC class entry

interface

Synopsis: *string*

The interface to which this class applies... Each interface must be listed only once.

mark

Synopsis: *unsigned short integer*

Mark that identifies traffic belonging to this class... This is an

unique integer between 1-255. Each class must have its own unique mark.

min-bandwidth

Synopsis: *string*

The minimum bandwidth this class should get,
when the traffic load rise... This can be
either a numeric value or a calculated expression
based on the bandwidth of the interface. A
fixed numerical value must only be a number -
its unit is specified in Minbw-unit.
A calculated expression is based on a fraction of the
'full' bandwidth, such as:
'full/3' for a third of the bandwidth and
'full*9/10' for nine tenths of the bandwidth.
In such a case, do not specify any Minbw-unit

minbw-unit

Synopsis: *string - one of the following keywords { none, bps, mbps, mbit, kbps, kbit }*

Default: *none*

(kbps, mbps...) Only if min-bandwidth is a **single numerical value**

max-bandwidth

Synopsis: *string*

The maximum bandwidth this class is allowed to use
when the link is idle... This can be either a numeric
value or a calculated expression based on the bandwidth
of the interface. A fixed numerical value must only be
a number - its unit is specified in Maxbw-unit.
A calculated expression is based on a fraction of the
'full' bandwidth, such as:
'full/3' for a third of the bandwidth and
'full*9/10' for nine tenths of the bandwidth.
In such a case, do not specify any Maxbw-unit

maxbw-unit

Synopsis: *string - one of the following keywords { none, bps, mbps, mbit, kbps, kbit }*

Default: *none*

(kbps, mbps...) Only if max-bandwidth is a **single numerical value**

priority

Synopsis: *unsigned short integer*

Default:

The priority in which classes will be serviced... Higher priority
classes will experience less delay since they are serviced
first. Priority values are serviced in ascending order
(e.g. 0 is higher priority than 1).

description

Synopsis: *string*

A description for this configuration item

Options

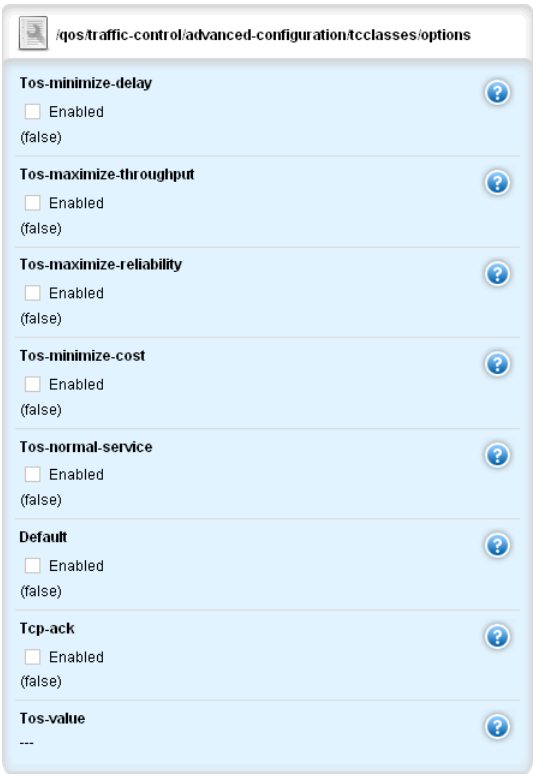

Figure 27.11. Options form

To display this form, navigate to *qos/traffic-control/advanced-configuration/tcclasses/{class}.*

IP Traffic matching with the ToS options take precedence over the mark rules.

tos-minimize-delay

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x10/0x10

tos-maximize-throughput

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x08/0x08

tos-maximize-reliability

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x04/0x04

tos-minimize-cost

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x02/0x02

tos-normal-service

Synopsis: *boolean*

Default: *false*

Value/mask encoding: 0x00/0x1e

default

Synopsis: *boolean*

Default: *false*

One default class per interface **must** be defined

tcp-ack

Synopsis: *boolean*

Default: *false*

All tcp ack packets into this class... This option should be

specified only once per interface.

tos-value

Synopsis: *string*

Custom classifier for the given value/mask...

The values are hexadecimal, prefixed by '0x'.

Ex.:

0x56[/0x0F]

The ToS-value field allows you to define a classifier for the given ToS value or value/mask combination of an IP packet's TOS byte. Value and Value/Mask are both specified in hexadecimal notation using the "0x" prefix. It is also possible to specify a diffserv marking, or DSCP (Diffserv Code Point). These are typically quoted as 6-bit values, and must be left-shifted (multiplied by 4) for use in the "tos=" field. For example, a DSCP of 0x2E (EF, or Expedited Forwarding) would be entered as 0xB8/0xFC (4 X 0x2E = 0xB8, and the two lowest order bits are masked by masking with 0xFC).

27.2.1.2.2. Devices

Devices characterize interfaces for traffic shaping, mostly for outbound bandwidth and the outgoing interface.

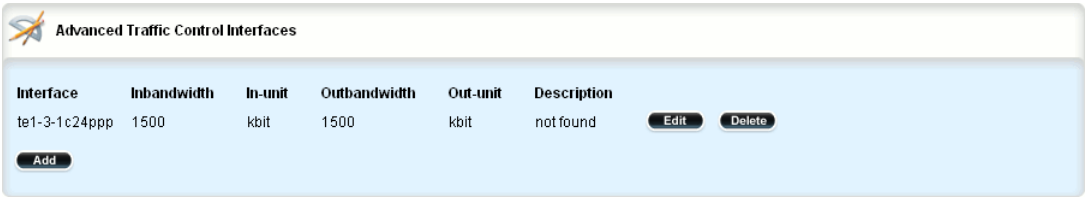

Figure 27.12. Advanced Traffic Control Interfaces table

The display this table, navigate to *qos/traffic-control/advanced-configuration/tcdevices*.

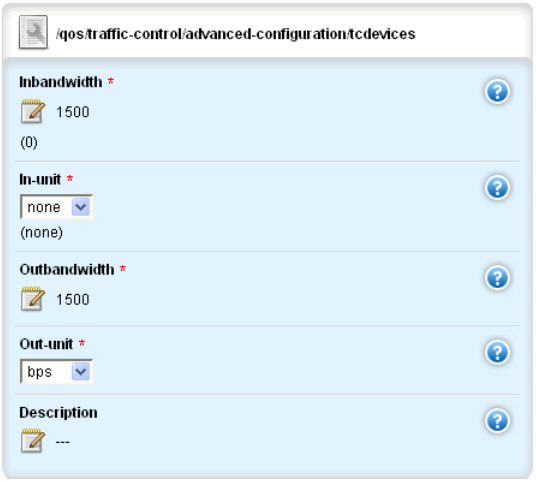

Figure 27.13. TC Devices form

The display this form, navigate to *qos/traffic-control/advanced-configuration/tcdevices/{interface}*.

interface

Synopsis: *string*

An interface to which traffic shaping will apply

inbandwidth

Synopsis: *unsigned short integer*

Default:

Incoming bandwidth - default: 0 = ignore ingress...

Defines the maximum traffic allowed for this interface in total,

if the rate is exceeded, the packets are dropped

in-unit

Synopsis: *string - one of the following keywords { none, bps, mbps, mbit, kbps, kbit }*

Default: *none*

Unit when incoming bandwidth is specified

outbandwidth

Synopsis: *unsigned short integer*

Maximum outgoing bandwidth... This is the maximum speed that can be

handled. Additional packets will be dropped. This is the

bandwidth that can be refrred-to as 'full' when defining classes.

out-unit

Synopsis: *string - one of the following keywords { bps, mbps, mbit, kbps, kbit }*

Outgoing bandwidth unit

description

Synopsis: *string*

A description for this configuration item

27.2.1.2.3. Rules

Rules define packet marking rules, usually for traffic shaping.

Figure 27.14. TCrules menu

The *tcrules* menu allows you to add, edit or remove a traffic classification rule. Add a new rule by selecting <Add tcrules>. Remove a tcrule by selecting next to a tcrule and click on an existing tcrule to modify it. Reorder rules by clicking next to the rule submenus.

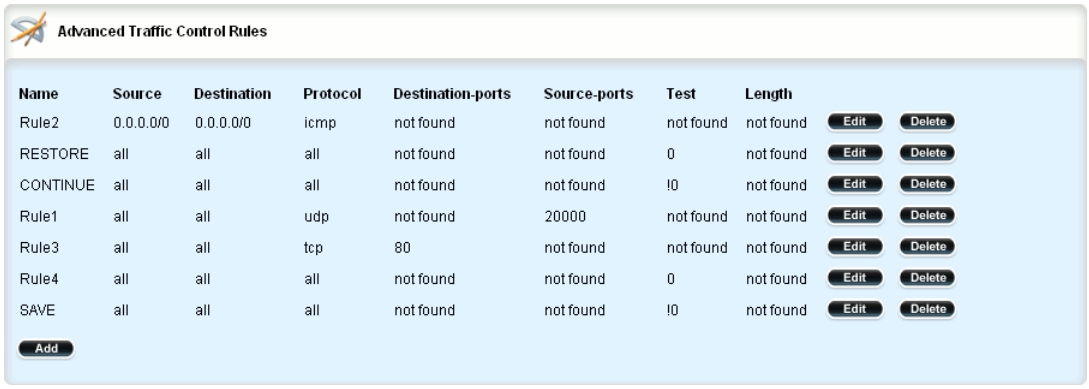
Advanced Traffic Control Rules
Name Source Destination Protocol Destination-ports Source-ports Test Length
Rule2 0.0.0.0/0 0.0.0.0/0 icmp not found not found not found not found Edit Delete
RESTORE all all all not found not found 0 not found Edit Delete
CONTINUE all all all not found not found !0 not found Edit Delete
Rule1 all all udp not found 20000 not found not found Edit Delete
Rule3 all all tcp 80 not found not found not found Edit Delete
Rule4 all all all not found not found 0 not found Edit Delete
SAVE all all all not found not found !0 not found Edit Delete
Add

Figure 27.15. Advanced Traffic Control Rules table

The display this table, navigate to *qos/traffic-control/advanced-configuration/tcrules*.

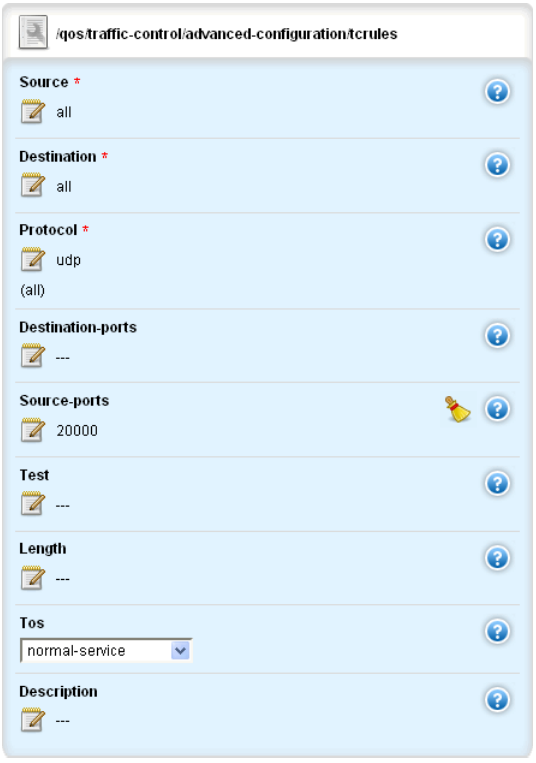

Figure 27.16. TCrules form

The display this form, navigate to *qos/traffic-control/advanced-configuration/tcrules/{rule}*.

name

Synopsis: *string*

A distinct name for this rule

source

Synopsis: *string*

IF name, comma-separated list of hosts or IPs, MAC addr, or 'all'...

When using MACs, use '~' as prefix and '-' as separator. Ex.:

~00-1a-6b-4a-72-34,~00-1a-6b-4a-71-42

destination

Synopsis: *string*

IF name, comma-separated list of hosts or IPs, or 'all'

protocol

Synopsis: *string*

Synopsis: *string - one of the following keywords { all, icmp, udp, tcp }*

Default: *all*

The protocol to match - default: all

destination-ports

Synopsis: *string*

(Optional) Comma-separated list of port names, port numbers or port ranges

source-ports

Synopsis: *string*

(Optional) Comma- separated list of port names, port numbers or port ranges

test

Synopsis: *string*

(Optional) Defines a test on the existing packet or connection mark...

Default is packet mark. For testing a connection mark, add

':C' at the end of the test value. Ex.:

Test if the packet mark is not zero:

!0

Test if the connection mark is not zero:

!0:C

length

Synopsis: *string*

(Optional) Match the length of a packet against a specific value or range of values...

Greater than and lesser than, as well as ranges are supported in the form of min:max. Ex.:

Equal to 64

64

Greater or equal to 65

65:

Lesser or equal to 65

:65

In-between 64 and 768

64:768

tos

Synopsis: *string*

Synopsis: *string - one of the following keywords { normal-service, minimize-cost, maximize-reliability, maximize-throughput, minimize-delay }*

(Optional) Type Of Service ...

A pre-defined TOS value or a numerical value. The

numerical value is hexadecimal. Ex.: 0x38

description

Synopsis: *string*

A description for this configuration item

Mark-choice

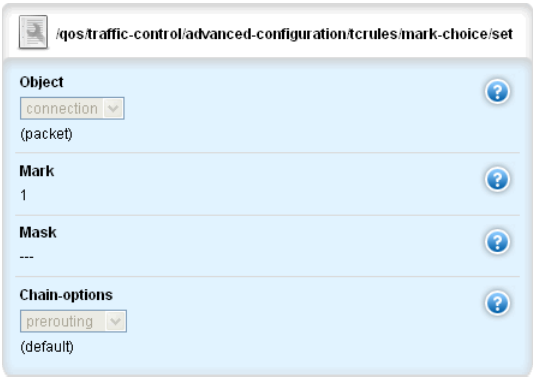

Figure 27.17. Set form

object

Synopsis: *string - one of the following keywords { connection, packet }*

Default: *packet*

Set the mark on either a packets or a connection

mark

Synopsis: *string*

Mark that corresponds to a class mark (decimal value)

mask

Synopsis: *string*

(optional) Mask to determine which mark bits will be set

chain-options

Synopsis: *string - one of the following keywords { prerouting, postrouting, forward }*

Default: *forward*

Chain where the set operation will take place

The chain-options field specifies the chain in which the rule will be processed.

- *Prerouting - Mark the connection in the PREROUTING chain.*

  This can be used with DNAT, SNAT and Masquerading rule in firewall. An example of such a rule is "Source.IP:192.168.2.101, Chain-option: preroute or default" but the actual Source.NAT address is 2.2.2.2.

- *Postrouting - Mark the connection in the POSTROUTING chain.*

  This can be used with DNAT, SNAT and Masquerading rules in the firewall. An example of such rule is "Destination.IP:192.168.3.101, Chain-option:preroute or default". In this case, the actual destination address is 192.168.3.101 but it will be translated to 192.168.3.33 by DNAT. Another example of a traffic control rule is ""Destination.IP:192.168.3.33, Chain-option:postrouting".

- *Forward - Mark the connection in the FORWARD chain.*

  This is the default chain option and it can be used for normal IP traffic without any address or port translation.

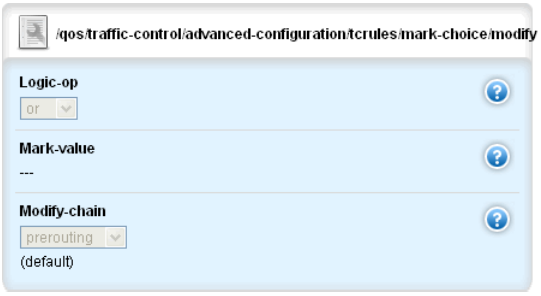

Figure 27.18. Modify form

logic-op

Synopsis: *string - one of the following keywords { or, and }*

Logical operation to perform on the current mark: AND/OR

mark-value

Synopsis: *string*

Mark to perform the operation with (decimal value)

modify-chain

Synopsis: *string - one of the following keywords { prerouting, postrouting, forward }*

Default: *forward*

Chain in which the operation will take place

Figure 27.19. Save form

value-mask

Synopsis: *string*

Mask to process the mark with

op-chain

Synopsis: *string - one of the following keywords { prerouting, forward }*

Default: *forward*

Chain in which the operation will take place

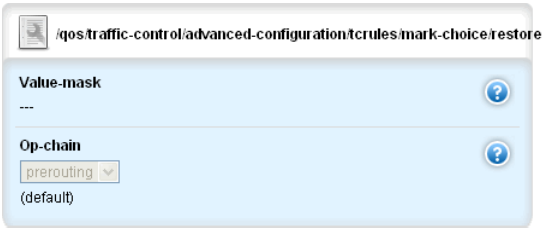

Figure 27.20. Restore form

value-mask

Synopsis: *string*

Mask to process the mark with

op-chain

Synopsis: *string - one of the following keywords { prerouting, forward }*

Default: *forward*

Chain in which the operation will take place

Figure 27.21. Continue form

continue-chain

Synopsis: *string - one of the following keywords { prerouting, forward }*

Default: *forward*

Chain in which the operation will take place

Hints on optimizing the TC Rule table

Every rule is processed in table order for every packet, unless a CONTINUE rule is matched, in which case processing stops. This can be used to improve efficiency in combination with the SAVE and RESTORE rules. For example, consider a TC Rules table organized roughly as follows (and in the same order):

- A RESTORE rule is used to restore the connection's mark to a matching, unmarked packet
- A CONTINUE if the mark is non zero
- Specific rules to check criteria to assign a mark, and finally
- A SAVE mark to connection if the mark is non zero (that is, a match was found above)

Using the above structure for the TC Rules table, only the first packet of any tcp or udp connection will have to go through all the rules, while every following packet will have its mark restored by the first rule, and then CONTINUE, skipping potentially many matching rules in the remainder of the table.

28. VRRP

28.1. VRRP Fundamentals

The Virtual Router Redundancy Protocol (VRRP) eliminates a single point of failure associated with statically routed networks by providing automatic failover using alternate routers. The RuggedRouter® VRRP daemon (keepalived) is an RFC 2338 version 2 compliant implementation of VRRP.

28.1.1. The Problem With Static Routing

Many network designs employ a statically configured default route in the network hosts. A static default route is simple to configure, requires little if any overhead to run and is supported by virtually every IP implementation. When dynamic host configuration protocol (DHCP) is employed, hosts may accept configuration for only a single default gateway.

Unfortunately, this approach creates a single point of failure. Loss of the router supplying the default route or the router's WAN connection results in isolating the hosts relying upon the default route.

There are a number of ways that may be used to provide redundant connections to the host. Some hosts can configure alternate gateways while others are intelligent enough to participate in dynamic routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First routing protocol (OSPF). Even when available, these approaches are not always practical due to administrative and operation overhead.

28.1.2. The VRRP Solution

VRRP solves the problem by allowing the establishment of a "virtual router group", composed of a number of routers that provide a specific default route. VRRP uses an election protocol to dynamically assign responsibility for the "virtual" router to one of the routers in the group. This router is called the VRRP Master. If the Master (or optionally its WAN connection) fails, the alternate (i.e. backup) routers in the group elect a new Master. The new master provides the virtual IP address and issues a gratuitous ARP to inform the network of where the gateway can be reached.

Because the host's default route does not change and MAC address is updated, packet loss at the hosts is limited to the amount of time required to elect a new router.

28.1.3. VRRP Terminology

Each physical router running VRRP is known as a VRRP Router. Two or more VRRP Routers can be configured to form a "Virtual Router". Each VRRP Router may participate in one or more Virtual Routers.

Each Virtual Router has a user-configured *Virtual Router Identifier* (VRID) and an *Virtual IP address* or set of IP addresses on the shared LAN. Hosts on the shared LAN are configured to use these addresses as the default gateway.

One router in the Virtual Router Group will be elected as the *Master*, all other routers in the group will be *Backups*.

Each router in the group will run at a specific *Priority*. The router with the highest priority is elected Master. The value of Priority varies from 1 to 255.

VRRP can also monitor a specified interface and give up control of a VRIP if that interface goes down.

In the following network, host 1 uses a gateway of 1.1.1.253 and host 2 uses a gateway of 1.1.1.252. The 1.1.1.253 gateway is provided by VRID 10. In normal practice router 1 will provide this virtual IP as its priority for VRID 10 is higher than that of router 2. If router 1 becomes inoperative or if its w1ppp link fails, it will relinquish control of VRIP 1.1.1.253 to router 2.

In a similar fashion host 2 can use the VRID 11 gateway address of 1.1.1.252 which will normally be supplied by router 2.

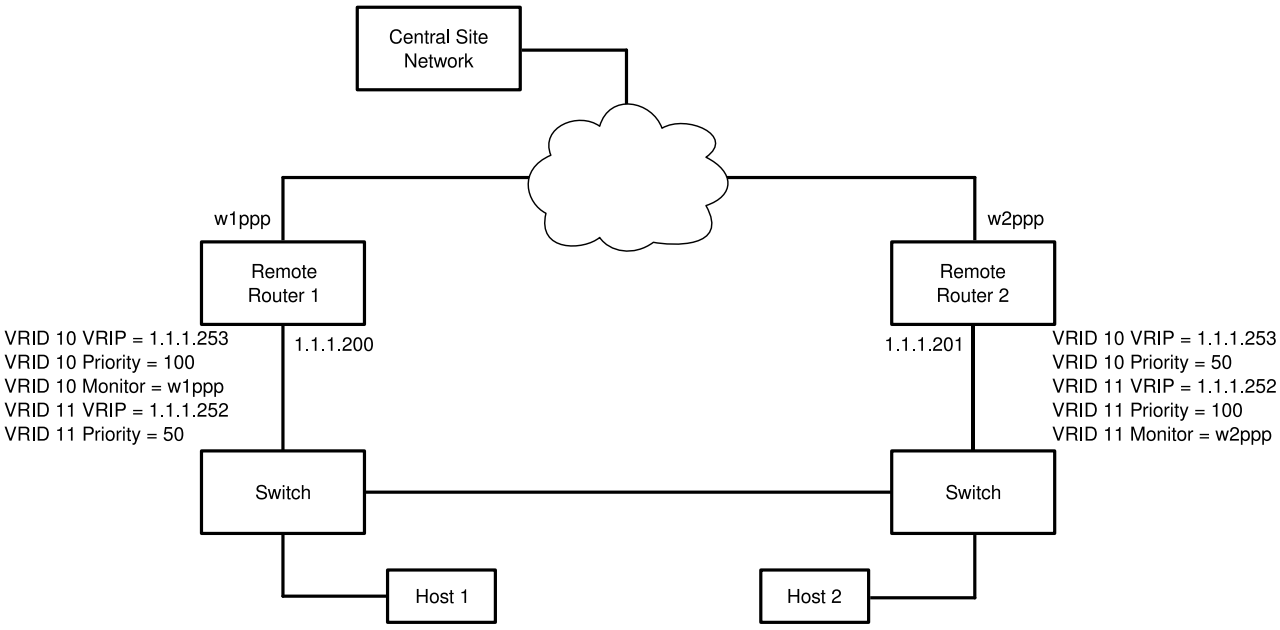

Figure 28.1. VRRP Example

In this example traffic from host1 will be sent through router 1 and traffic from host2 through router 2. A failure of either router (or its wan link) will be recovered by the other router.

Note that both routers can always be reached by the hosts at their “real” IP addresses.

Two or more VRRP instances can be assigned to be in the same *VRRP Group*, in which case, they can fail over together.

In the following network, both host 1 and host 2 use a gateway of 192.168.3.10. The external side can access the internal side by gateway 192.168.2.10. The VRID_20 and VRID_21 are grouped together. Normally the Router 1 will provide both internal and external access gateway as its priority is higher than those on Router 2. When either internal or external side of Router 1 becomes inoperative, it will remove the control of both VRIP 192.168.2.10 and 192.168.3.10 and gives the control to Router 2.

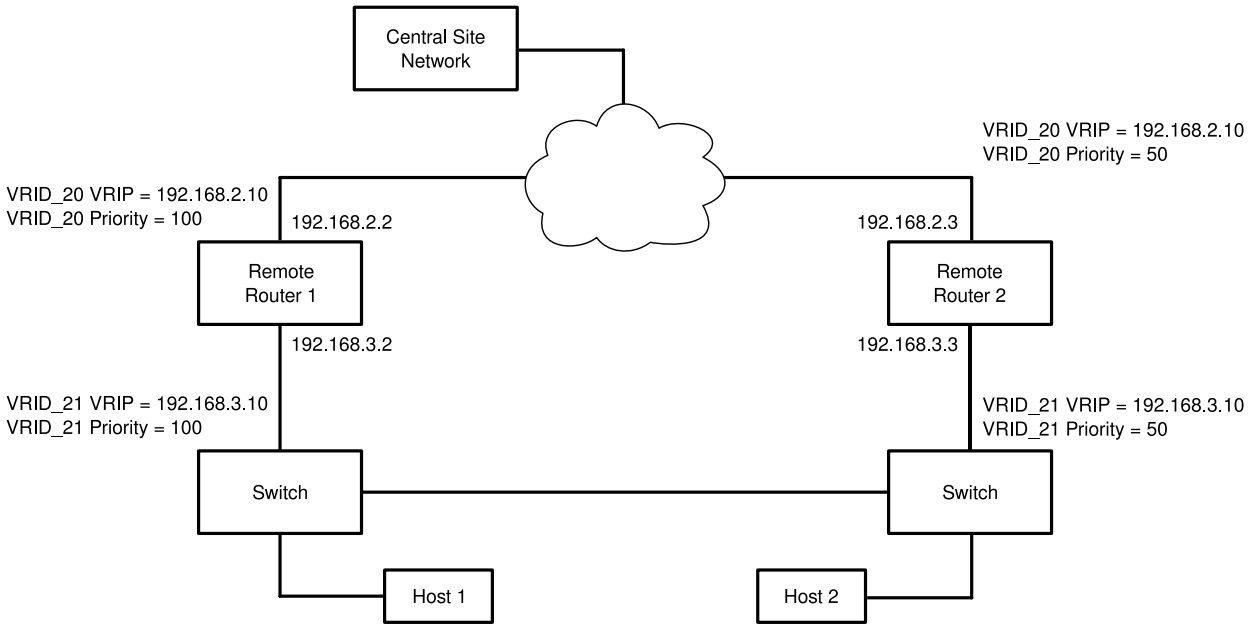

Figure 28.2. VRRP Group Example

Other VRRP parameters are the *Advertisement Interval* and *Gratuitous ARP* Delay.

The advertisement interval is the time between which advertisements are sent. A backup router will assume mastership *four advertisement intervals* after the master fails, so the minimum fail-over time is four seconds. If a monitored interface goes down, a master router will immediately signal an election and allow a backup router to assume mastership.

The router issues a set of gratuitous ARPs when moving between master and backup state. These unsolicited ARPs teach the hosts and switches in the network of the current MAC address and port associated with the VRIP. The router will issue a second set of ARPs after the time specified by the Gratuitous ARP delay.

28.2. VRRP Configuration

Figure 28.3. VRRP Menu

The VRRP menu is accessible from the main menu under *services* at *services/vrrp*.

Figure 28.4. Virtual Router Redundancy Protocol (VRRP) Form

The Virtual Router Redundancy Protocol (VRRP) form appears on the same screen as the *VRRP* menu. In the Virtual Router Redundancy Protocol (VRRP) form, enable or disable the VRRP service.

Enable VRRP Service

Enables VRRP Service.

Router ID

Synopsis: *string*

The router ID for VRRP logs.

Figure 28.5. VRRP Group Table

The VRRP Group table displays configured VRRP groups. To display this table, navigate to *services/vrrp/group*.

Group Name

Synopsis: *string*

The group name.

Figure 28.6. VRRP Instance Table

To display this table, navigate to *services/vrrp/instance*.

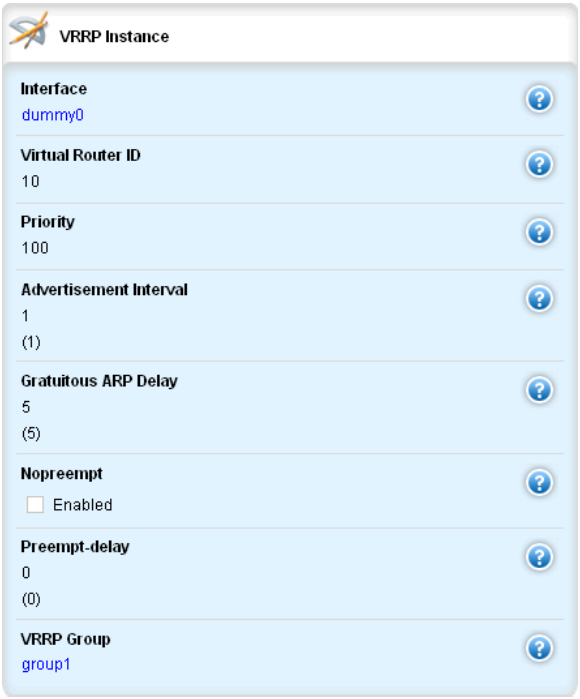

Figure 28.7. VRRP Instance Form

The VRRP Instance Form is used when configuring a VRRP instance. To display this form, navigate to *services/vrrp/instance/VRID20*.

Instance Name

Synopsis: *string*

The VRRP instance name.

Interface

Synopsis: *A string*

The interface that VRRP packets are sent on.

Virtual Router ID

Synopsis: *unsigned byte*

The Virtual Router ID. All routers supplying the same VRIP should have the same VRID.

Priority

Synopsis: *unsigned byte*

The priority of VRRP instance. For electing MASTER, highest priority wins.

Advertisement Interval

Synopsis: *unsigned byte*

Default: *1*

The advertisement interval (in seconds).

Gratuitous ARP Delay

Synopsis: *unsigned byte*

Default: *5*

Gratuitous ARP delay (in seconds). This controls the number of seconds after the router changes between the master and the backup state before a second set of gratuitous ARPs are sent.

nopreempt

Allows lower priority machine to maintain master role, even when a higher priority machine comes back online.

preempt-delay

Synopsis: *unsigned integer*

Default:

Seconds after startup until preemption.

use-virtual-mac

Use virtual MAC.

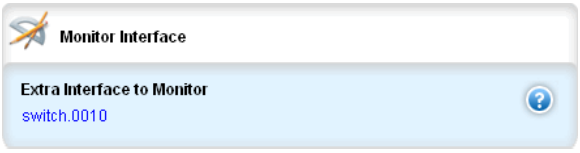

Figure 28.8. Monitor Interface Form

To display this form, navigate to *services/vrrp/instance/VRID20/monitor*.

An Extra Interface to Monitor causes VRRP to release control of the VRIP if the specified interface stops running.

Extra Interface to Monitor

Synopsis: *A string*

The interface name.

Figure 28.9. VRIP IP Address Table

The VRIP IP Address Table shows configured VRIP IP addresses associated with a VRID. To display this table, navigate to *services/vrrp/instance/VRID20/vrip*.

Virtual IP Address/Netmask

Synopsis: *IPv4 address and prefix in CIDR notation*

Virtual IP Address/Netmask.

28.2.1. VRRP Status

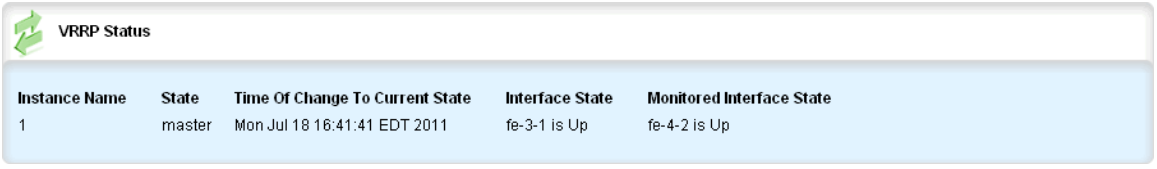

Figure 28.10. VRRP Status Table

To display this table, navigate to *services/vrrp/status*.

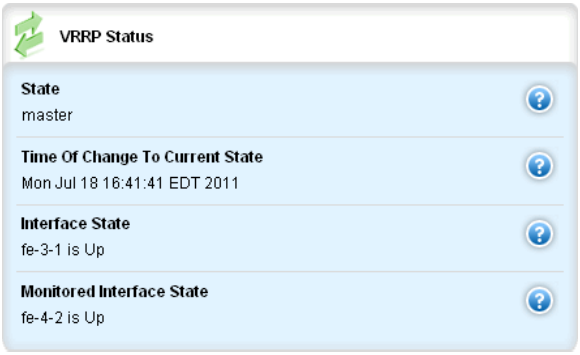

Figure 28.11. VRRP Status Form

To display this form, navigate to *services/vrrp/status/{number}*.

Instance Name

Synopsis: *string*

The VRRP instance name.

State

Synopsis: *string*

The VRRP instance state.

Time Of Change To Current State

Synopsis: *string*

The time of change to the current state.

Interface State

Synopsis: *string*

The VRRP interface state.

Monitored Interface State

Synopsis: *string*

Monitors the interface state.

29. Link Failover

Link failover provides an easily configured means of raising a backup link upon the failure of a designated main link. The main and backup links can be Ethernet, Cellular or Dial Modem, T1/E1, T3/E3, or DDS.

Link failover can back up to multiple remote locations, managing multiple main-to-backup link relationships. When the backup link is a modem, many “profiles” of dialed numbers can exist, each serving as a distinct backup link.

Link failover can back up a permanent, high-speed WAN link to a permanent, low-speed WAN link. Use this function when OSPF cannot be employed, such as on public links.

You can also use link failover to migrate the default route from the main link to the backup link.

The time after a main link failure to backup link startup, and the time after a main link recovery to backup link stoppage, are configurable. The link failover function also provides failover status information and a test of the failover settings.

29.1. Path Failure Discovery

To discover a primary path failure (in this example, through Network A), the link backup daemon inspects the link status of the main link and sends a regular ping to a designated host (or to a dummy address on the router). In this way, network link failures can be discovered. It is essential that the designated host or that the dummy address always respond to the ping.

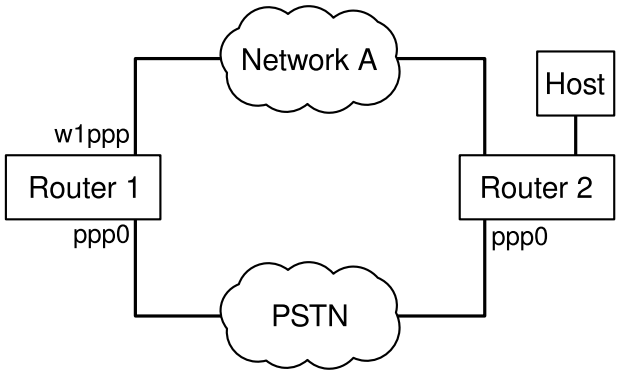

Figure 29.1. Link Backup Example

The daemon considers the main link to have failed, even if its link status is “up”, if the remote host fails to respond to a configurable number of pings after waiting a configurable timeout for each ping.

29.2. Using Routing Protocols and the Default Route

If the main trunk is on a private network, use a routing protocol to ensure that an alternate route to the end network is learned after the backup trunk is raised. Ensure that OSPF/RIP are configured to operate on the secondary trunk, assigning the secondary trunk a higher metric cost than that of the main trunk.

If the main trunk is on a public network, use the “transfer default route” feature.

The default route of the main trunk to be backed up using “transfer default route” must be assigned a metric greater than one.

29.3. Configuring Link Failover

To view the list of configured link failover settings, navigate to */services/link-failover*.

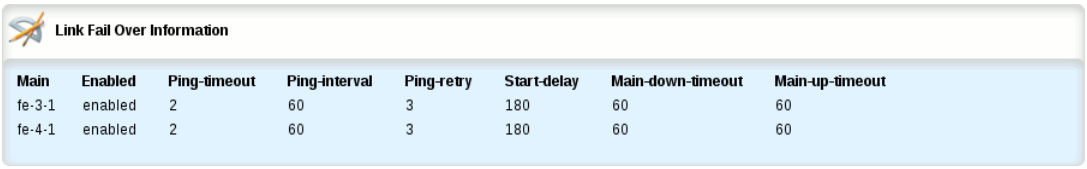

Link Fail Over Information

Main	Enabled	Ping-timeout	Ping-interval	Ping-retry	Start-delay	Main-down-timeout	Main-up-timeout
fe-3-1	enabled	2	60	3	180	60	60
fe-4-1	enabled	2	60	3	180	60	60

Figure 29.2. Link Fail Over Information Table

To configure link failover, do the following:

- set the link failover settings. See Section 29.3.1, “Configuring the Link Failover Settings”.
- add a link failover backup interface. See Section 29.3.2, “Setting a Link Failover Backup Interface”.
- add a link failover ping target. See Section 29.3.3, “Setting a Link Failover Ping Target”.
- enable link backup on demand. See Section 29.3.4, “Link Backup On Demand”.

After configuring link failover, you can do the following:

- view the link failover status. See Section 29.3.5, “Viewing Link Failover Status”.
- view the link failover log. See Section 29.3.6, “Viewing the Link Failover Log”.
- test the link failover settings. See Section 29.3.7, “Testing Link Failover”.

29.3.1. Configuring the Link Failover Settings

To configure the link failover settings:

- In edit mode, navigate to */services/link-failover* and click **<Add link-failover>**.
- On the *Key settings* form, select an interface from the list and click **Add**.
- On the *Link Fail Over Settings* form, set the link failover parameters.
- Commit the changes.

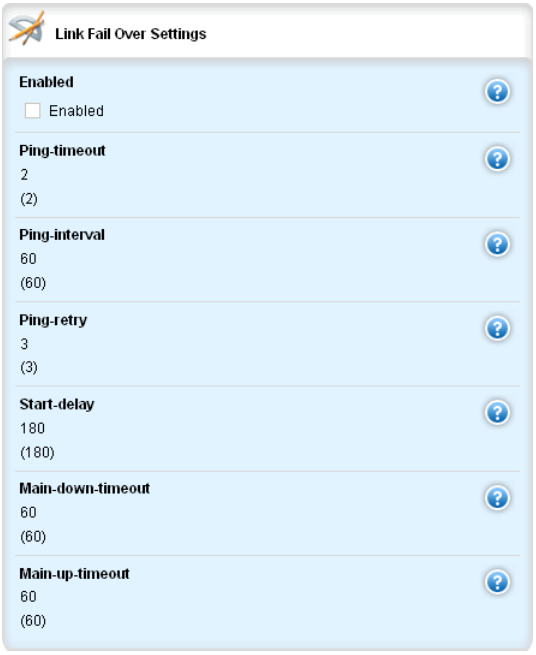

Figure 29.3. Link Fail Over Settings form

enabled

Enables this link backup.

ping-timeout

Synopsis: *integer*

Default: *2*

The time interval, in seconds, before immediately retrying a ping.

ping-interval

Synopsis: *integer*

Default: *60*

The time interval, in seconds, between ping tests.

ping-retry

Synopsis: *integer*

Default: *3*

The number of ping retries before construing a path failure.

start-delay

Synopsis: *integer*

Default: *180*

The delay time, in seconds, when first starting link failover.

main-down-timeout

Synopsis: *integer*

Default: *60*

The delay time, in seconds, that the main trunk is down before starting the backup trunk.

main-up-timeout

Synopsis: *integer*

Default: *60*

The delay time, in seconds, to confirm that the main trunk is up (returned to service) before stopping the backup trunk.

29.3.2. Setting a Link Failover Backup Interface

A backup interface is the interface to which link failover switches when the main interface is determined to be down. You can add up to three backup interfaces to each link failover configuration.

To set a link failover backup interface:

- In edit mode, navigate to */services/link-failover{link failover ID}/backup* and click **<Add backup>**.
- On the *Key settings* form, select an interface from the list and click **Add**.
- On the *Backup Settings* form, set the backup parameters.
- Commit the changes.

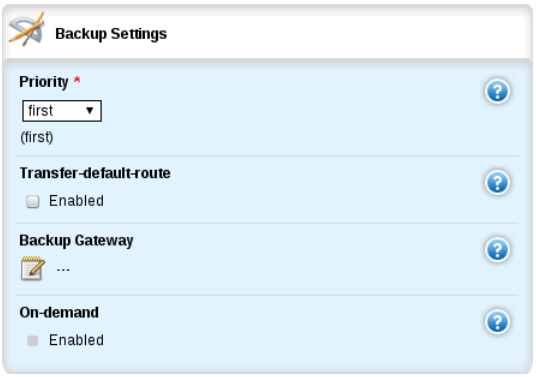

Figure 29.4. Backup Settings form

priority

Synopsis: *string - one of the following keywords { first, second, third }*

Default: *first*

The priority which is applied to the backup interface when switching

transfer-default-route

Transfer default gateway on switching main and backup interface.

The default route on the main interface must have a 'distance' greater than one.

Backup Gateway

Synopsis: *A string conforming to: "(([0-9]\|[1-9][0-9]\|1[0-9]{2}\|2[0-4][0-9]\|25[0-5])\.){3}([0-9]\|[1-9][0-9]\|1[0-9]{2}\|2[0-4][0-9]\|25[0-5])"*

The IP address of the backup gateway

on-demand

Synopsis: *boolean*

Displays the status of the interface's On-demand option. When enabled, link failover can set the interface up or down as needed; the interface is down until needed by link failover. When disabled, link failover cannot set the interface up or down; by default, the interface is always up.

29.3.3. Setting a Link Failover Ping Target

A link failover ping target is an IP address that link failover pings to determine if the main link is down. The address can be a dedicated host or a dummy address on a router. You can add up to three link failover ping targets to each link failover configuration.

To set a link failover ping target:

- In edit mode, navigate to */services/link-failover{link failover ID}/backup* and click **<Add target>**.
- On the *Key settings* form, enter the IP address of the host to ping and click **Add**.
- Commit the changes.

> *Link failover pings each target separately. If all targets are down, the main link is considered to be down and it fails over to the backup interface. Backup links are used in the order of their Priority setting (first, second, and then third), always starting with the first priority interface. When a higher-priority interface becomes available again, the system reverts to the higher priority interface. For example, if the second priority interface is active, the system switches back to the first priority interface when the first priority interface becomes available again.*

29.3.4. Link Backup On Demand

Use the *On-demand* option to keep interfaces down until they are needed by link failover:

- When the *On-demand* option is enabled on an interface, the interface is down by default. The interface is brought up when needed by the link failover function, and is brought down again when no longer needed.
- When the *On-demand* option is not enabled on an interface, the interface is always up by default.

The *On-demand* option can be set for the following interfaces:

- eth: on the Routable Ethernet Ports form at */interface/eth{interface id}*
- wan mlppp connections: on the Multilink PPP form at */interface/wan{interface id}/t1/channel{channel id}/connection/mlppp*
- wan ppp connections: on the PPP form at */interface/wan{interface id}/t1/channel{channel id}/connection/ppp*
- wan framerelay connections: on the DLCI form at */interface/wan{interface id}/t1/channel{channel id}/connection/framerelay/dlci*

After enabling the on-demand option, the On-demand field indicates the option's status on the Backup Settings form.

29.3.5. Viewing Link Failover Status

The *Link Fail Over Status* form displays the current link failover status. To view the link failover status in normal or edit mode, navigate to */services/link-failover{interface id}/status*.

Figure 29.5. Link Fail Over Status form

main-link-status

Synopsis: *string*

The main link status.

backup-link-status

Synopsis: *string*

The backup link status.

main-ping-test

Synopsis: *string*

Results of the pinging target using the main interface.

time-of-last-state-change

Synopsis: *string*

The time of the last state change.

link-backup-state

Synopsis: *string*

The backup link state.

backup-interface-in-use

Synopsis: *string*

The name of the backup interface that is being used.

29.3.6. Viewing the Link Failover Log

The *Link Fail Over Logs* form displays a log of link failover events. To view the link failover log in normal or edit mode, navigate to */services/link-failover{interface id}/log* and click **Perform**.

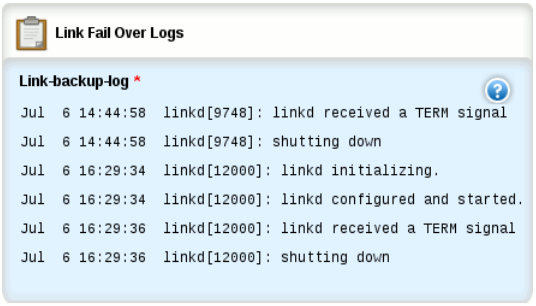
Link Fail Over Logs
Link-backup-log *
Jul 6 14:44:58 linkd[9748]: linkd received a TERM signal
Jul 6 14:44:58 linkd[9748]: shutting down
Jul 6 16:29:34 linkd[12000]: linkd initializing.
Jul 6 16:29:34 linkd[12000]: linkd configured and started.
Jul 6 16:29:36 linkd[12000]: linkd received a TERM signal
Jul 6 16:29:36 linkd[12000]: shutting down

Figure 29.6. Link Fail Over Logs form

29.3.7. Testing Link Failover

You can test your link failover settings to confirm that each link failover configuration works properly. To launch the test, you specify for how long the system should operate on the backup interface, and for how long the system should delay before starting the test. You can also cancel a test while it is in progress. Cancelling the test returns the interfaces to their pre-test condition.

While the test is running, monitor the *Link Fail Over Status* form to observe the main and backup link status, ping test results, state change, backup state, and backup interface information. As the test progresses, this information changes as link failover switches from the main interface to the backup interface. For more information on the *Link Fail Over Status* form, see Section 29.3.5, “Viewing Link Failover Status”.

To launch a link failover test:

- In normal mode or edit mode, navigate to */services/link-failover{interface id}/start-test*.

- On the *Link Fail Over Test Settings* form, set the test parameters.
- On the *Trigger Action* form, click **Perform**.

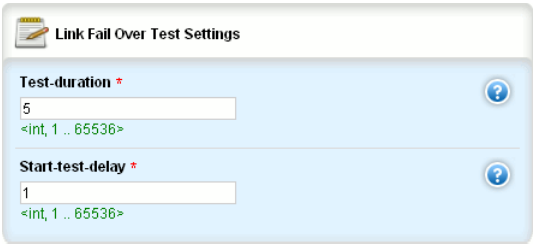

Figure 29.7. Link Fail Over Test Settings form

Test-duration

The amount of time (in minutes) to run the test before restoring service to the main trunk.

Start-test-delay

The amount of waiting time (in minutes) before running the test.

Part IV. Appendices

Upgrading Software	Appendix A, *Upgrading Software*
RADIUS Server Configuration	Appendix B, *RADIUS Server Configuration*
Setting Up An Upgrade Server	Appendix C, *Setting Up An Upgrade Server*
GNU General Public License	Appendix D, *GNU General Public License*

Appendix A. Upgrading Software

To launch a ROX™ operating system software upgrade, follow the procedure outlined below.

A.1. Preparing The Software Upgrade

The first step in a ROX™ software upgrade is to configure the location of the software upgrade repository and the version of software to which to upgrade.

At the top of the screen, click **Edit Private** to access the Edit Private view. The screen in Edit Private view is shown in the figure below.

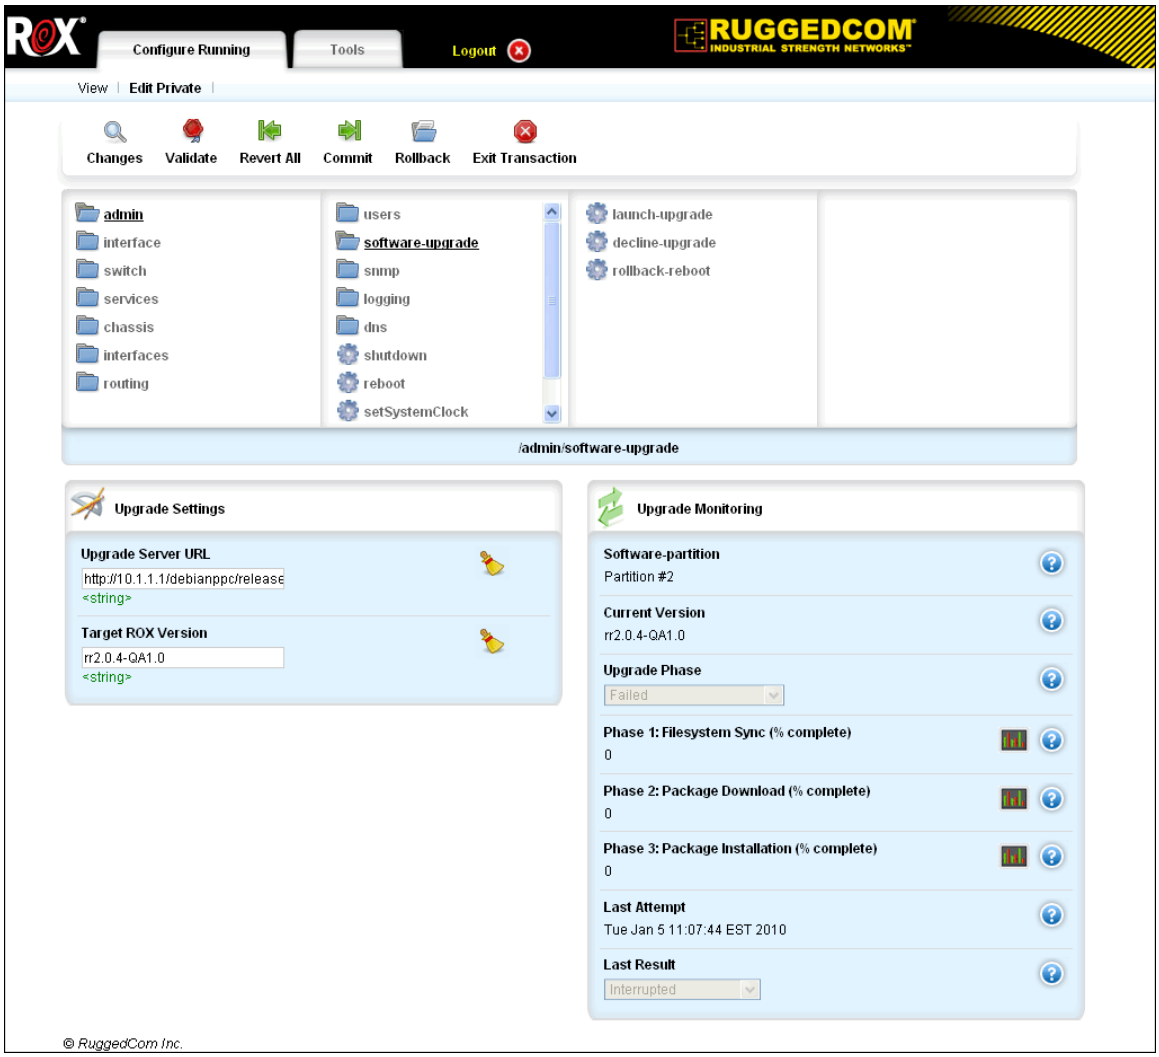

Figure A.1. The Software Upgrade Menu Interface

In the Upgrade Settings form, enter the location of the software upgrade server Upgrade Server URL field, and the version string of the desired version of ROX™ in the Target ROX Version field.

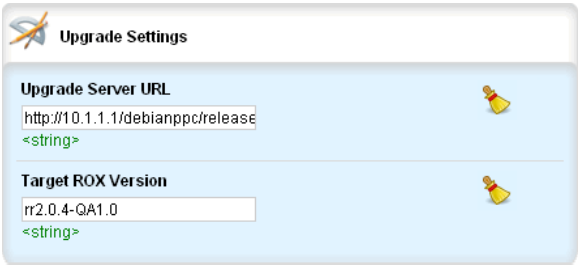

Figure A.2. Entry Fields in Upgrade Settings Form

After completing the information in the Upgrade Settings form, click the **Commit** button ( ) at the top of the screen. A dialog box will appear, prompting you to commit your changes. Click the **OK** button.

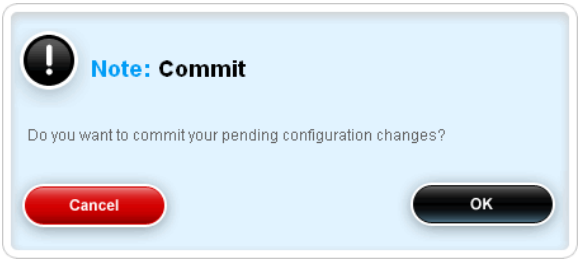

Figure A.3. Pending Commit

A dialog box will appear, informing you that the configuration has been committed. Click the **OK** button.

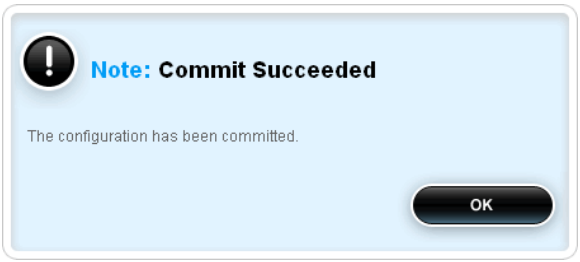

Figure A.4. Commit Succeeded

A.2. Launching The Upgrade

Click on the *launch-upgrade* menu action. Then, click the **Perform** button on the Launch Upgrade form. A dialog box will assert that: "Configurations will be locked until next boot." Click the **OK** button.

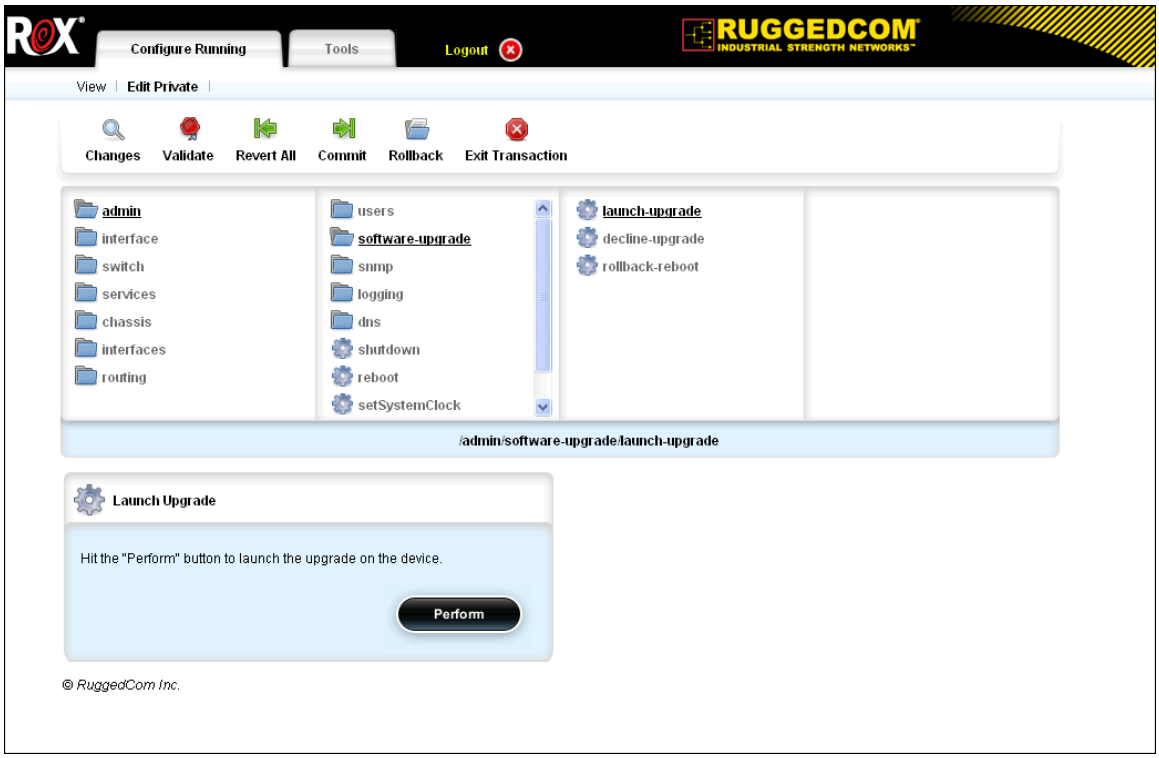

Figure A.5. Launch Upgrade

The Success! and Upgrade Options messages shown below indicate that the upgrade has been launched.

Success!
The upgrade has been performed.
Upgrade Options
Launch-result *
The upgrade has been launched. View 'upgrade-progress' to monitor its status.

Figure A.6. Upgrade Launched Dialogs

Click the **Exit Transaction** ( ) button at the top of the screen to return to the *View* mode.

A.3. Monitoring The Software Upgrade

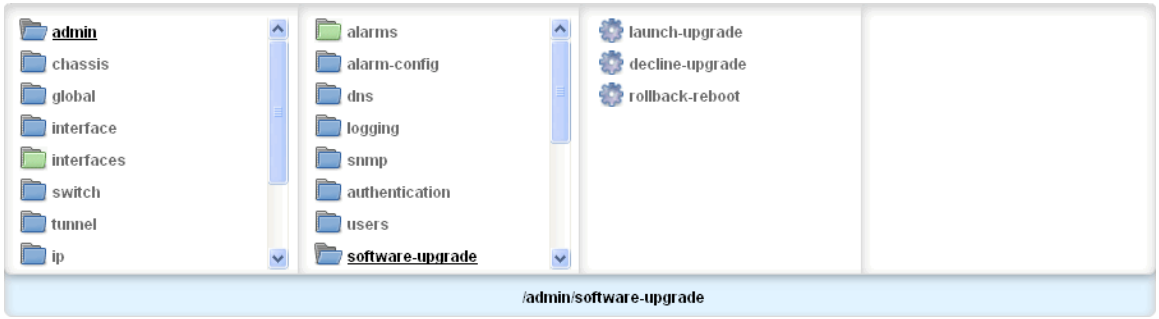

Figure A.7. Software-Upgrade Menu

Click on the Software-Upgrade menu to view the Upgrade Monitoring form. The Upgrade Monitoring form shows the real-time progress of the Upgrade procedure. The software upgrade progresses through four phases:

- Estimating upgrade size
- Copying filesystem
- Downloading packages
- Installing packages

The final reported status is either *Completed successfully* or, if an error prevented the completion of any one of the above four phases, *Failed*. These phases are shown in real time in the *Upgrade Phase* field on the Upgrade Monitoring Form, below.

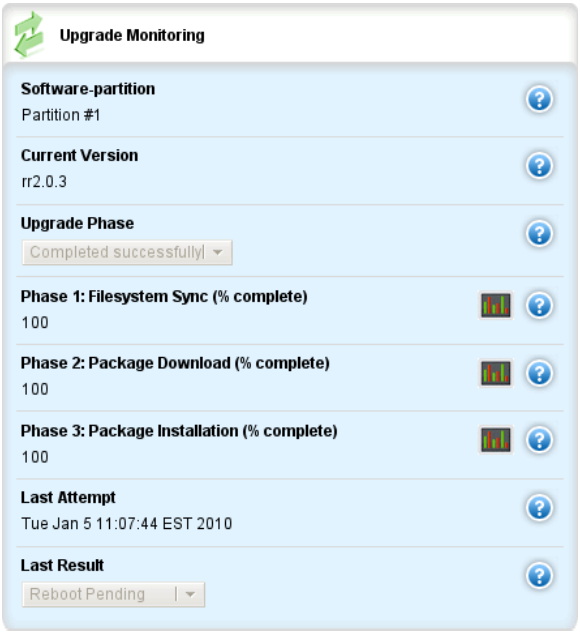

Figure A.8. Upgrade Monitoring Form in Reboot-pending Stage

Once the upgrade has successfully installed, the three phase progress indicator fields indicate 100% complete. In this state, the Upgrade Monitoring form will display *Reboot Pending* in the *Last Result* field. This indicates that a system reboot is required in order for the newly installed software upgrade to take effect.

Reboot the system.

When system has been rebooted to run the newly upgraded software installation, the *Current Version* field will reflect the ROX version number:

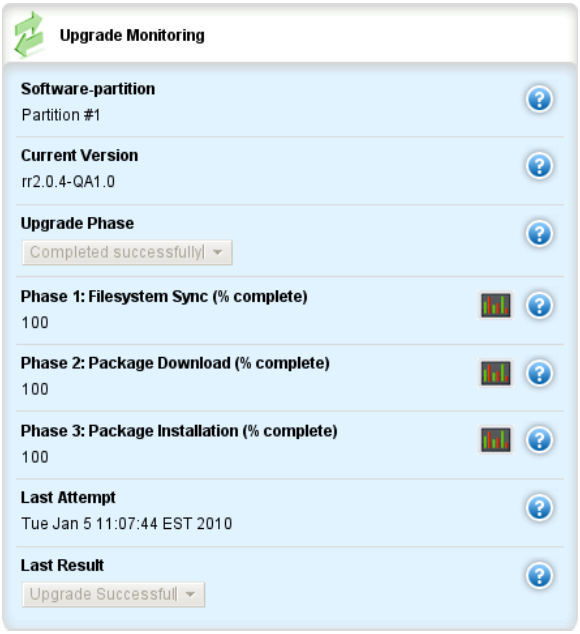

Figure A.9. Upgrade Monitoring Form Showing Successful Upgrade

software-partition

synopsis: a string of at most 31 characters

The current active partition number. The unit has two software partitions: #1 and #2. Upgrades are always performed to the other partition.

current-version

synopsis: a string of at most 31 characters

The current operating software version.

upgrade-state

synopsis: string - one of { Failed, Completed successfully, Unknown state, Installing packages, Downloading packages, Copying filesystem, Estimating upgrade size, Inactive }

The current phase or state of the upgrade.

filesystem-copy

synopsis: integer in the range [0 to 100]

Phase 1 of the upgrade involves synchronizing the filesystem with the partition to which we are upgrading.

This reflects the estimated percent complete.

package-download

synopsis: integer in the range [0 to 100]

Phase 2 of the upgrade downloads all packages that require an update. This reflects the estimated percent complete.

package-installation

synopsis: integer in the range [0 to 100]

Phase 3 of the upgrade installs all packages that require an update. This reflects the estimated percent complete.

last-upgrade-attempt

synopsis: a string of at most 64 characters

The date and time of completion of the last upgrade attempt.

last-upgrade-result

synopsis: string - one of { Interrupted, Declined, Not Applicable, Reboot Pending, Unknown, Upgrade Failed, Upgrade Successful }

Indicates whether or not the last upgrade completed successfully

Appendix B. RADIUS Server Configuration

This section describes the configuration procedures for two popular RADIUS servers, "FreeRADIUS" and the Microsoft Windows "Internet Authentication Service" in order to create and manage accounts that are able to access resources on RuggedRouter®. There are four RADIUS attributes required for the configuration of accounts to access services on RuggedRouter®. The following table shows the RADIUS attributes required by RuggedRouter® for accounts that are designated to use one or more of the "login", "ppp", or "ssh" services:

RADIUS Attribute	login	ppp	ssh
User ID	required	required	required
Password	required	required	required
NAS-Identifier			
RuggedCom-Privilege-level			

Table B.1. Required Attributes for various RADIUS services

Every account to be authenticated on behalf of the RuggedRouter® must have a user ID and password. The RADIUS "NAS-Identifier" attribute may optionally be used to restrict which service an account may access:

- *login*
- *ppp*
- *ssh*

Accounts that do not specify a "NAS-Identifier" attribute may access any RuggedRouter® service upon authentication. Accounts may also be defined to have access to one or several services. For more information on these services on RuggedRouter®, please refer to RADIUS, ROXII, and Services.

You must all the following information to the vendor-specific extensions of the chosen RADIUS server:

- RuggedCom uses Vendor number 15004.
- "RuggedCom-Privilege-level" is attribute 2, of type "string".
- "RuggedCom-Privilege-level" must take one of the following three values:
  - "admin"
  - "operator"
  - "guest"

B.1. PPP / CHAP and Windows IAS

In order for Windows IAS to authenticate PPP connections that use the CHAP authentication protocol, IAS must be made to store passwords using what it calls "reversible encryption".

1. Ensure that CHAP authentication is enabled in the Remote Access Policy.
2. In the Active Directory settings for each PPP user, select "Store password using reversible encryption".

Appendix C. Setting Up An Upgrade Server

The RuggedCom software upgrade mechanism requires a repository of software to be available. The following instructions detail:

- Requirements for a repository server,
- Initial set up of a repository,
- Upgrading the repository to the latest release,
- Maintain separate releases streams for different groups of routers,
- Setting up one router to test new releases
- Configuring the network routers.

C.1. Upgrade Server Requirements

In order to establish a repository you will need a host that is accessible to the units that will be upgraded. This host must be able to act as a web server or ftp server. The host must also be able to access the RuggedCom web site in order to download new releases of software from RuggedCom.

The server requirements are fairly modest. The principal requirements are for disk space, bandwidth and the ability to serve an adequate number of http sessions.

Each software release will require approximately 75 Mb of disk space. Note that this figure includes an entire software image, most upgrades will involve the transfer of only a small fraction of this amount. A large number of such releases could easily be stored on a system of only modest capabilities. In practice, only one or two releases are usually all that need be kept.

The bandwidth requirements are determined by the many factors including the number of units, size of upgrade, when the routers upgrade, each unit's upgrade is bandwidth limited to 500kbps by default. Most web servers can serve files to the limit of the network interface bandwidth, so even a modest (e.g. 486 class machine) would prove acceptable.

The server should be able to accept at least as many http or ftp connections as there are upgradable units in the network. In practice you will configure the units to have staggered upgrade times in order to minimize the impact of upgrading on the network. A large upgrade (or a low bandwidth limiting value at each router) may cause all the routers to be upgrading at any one time.

C.2. Initial Upgrade Server Setup

You must create a directory on the web server to hold the releases for the router. The directory can have any name, such as “ruggedBackbone”.

Some administrators like to test the upgrade to the new release before making it available at the repository-url to which all the routers on their network are pointing. Perhaps this is desired if you have automated the routers to run the upgrade periodically, as launching an upgrade to a repository with the same release returns with no action. In this case simply create a separate directory in the webserver such as "ruggedBackboneTest", so the new release is not available to the routers on your network.

These directory names will be used in examples in the remainder of this section.

Ensure that the web server publishes these directories.

C.3. Upgrading The Repository

Releases are obtained from the RuggedCom web site as ZIP files. Download the ZIP file to your regular and/or test release directories and unzip them. You may delete the original ZIP file if desired.

The ZIP file name will be in the form rrX.Y.Z.zip. The major release number ‘X’ is changed when major new functionality (often hardware related) is offered. The minor release number ‘Y’ is increased when new features are added or serious bugs fixed, and the patch release number ‘Z’ is increased when minor bug repairs are made.

The zip file will extract to a directory that has the same name as the major release, e.g. “rr2”. As subsequent releases are made, they will also be extracted into this directory.

C.4. Setting Up The Routers

Suppose you have just unzipped rr2.1.1.zip into “ruggedBackbone” (or "ruggedBackboneTest" if you do not want your routers to point to it yet) on a server available to the network at server.xyz.net.

There are now two approaches available to you for the upgrade settings on the unit.

The first method described allows you to configure the unit to always upgrade to the last release that was installed on the server. The advantage of this approach is that you do not have to update upgrade configurations each time you obtain a new release. On the RuggedBackbone configure the *repository-url* parameter to “http://server.xyz.net/ruggedBackbone/rr2” and the *target-release* parameter to “current”. You can proceed to launch the upgrade on the ruggedBackbone (this can be done via CLI, web user interface, and NETCONF - see the appropriate user guide for details).

The second method allows you to configure the target release version explicitly. Some administrators may prefer this approach for sake of clarity, but this has the disadvantage of having to update all the units with the release version each upgrade. On the RuggedBackbone configure the *repository-url* parameter to “http://server.xyz.net/ruggedBackbone/rr2” and the *target-release* parameter to "rr2.1.1".

C.5. Using Microsoft Internet Information Services (IIS) Manager 6.0 or Higher as a ROX Upgrade Repository

When using Microsoft Internet Information Services (IIS) Manager 6.0 or higher as your ROX upgrade repository, you must add a new application/octet-stream MIME type named ***** to the IIS properties. The new MIME type is required for IIS to consider ROX upgrade packets as an application/octet stream. If the new MIME type is not added, ROX upgrades will fail.

To add the new MIME type, perform the following steps on your IIS server.

Procedure C.1. Adding the * MIME Type

1. In the Windows **Start** menu, right-click on **My Computer** and select **Manage**. The *Computer Management* dialog appears.
2. Under **Services and Applications**, locate the **Internet Information Services (IIS) Manager** node. Right-click on your ROX upgrade repository website and select **Properties**. The *Properties* dialog appears.
3. Select the **HTTP Headers** tab and click **MIME Types**. The *MIME Types* dialog appears.
4. Click **New**. The *MIME Type* dialog appears.
5. In the *Extension* field, type *****.

   In the *MIME type* field, type **application/octet-stream**.
6. Click **OK** on the *MIME Type*, *MIME Types*, and *Properties* dialog boxes.

Appendix D. GNU General Public License

Version 2, June 1991
Copyright © 1989, 1991 Free Software Foundation, Inc.

Free Software Foundation, Inc.
51 Franklin Street, Fifth Floor,
Boston, MA 02110-1301
USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.
Version 2, June 1991

D.1. Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software - to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps:

1. copyright the software, and
2. offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

D.2. TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

D.2.1. Section 0

This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The “Program”, below, refers to any such program or work, and a “work based on the Program” means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term “modification”.) Each licensee is addressed as “you”.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

D.2.2. Section 1

You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

D.2.3. Section 2

You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: If the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

D.2.4. Section 3

You may copy and distribute the Program (or a work based on it, under Section 2 in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

D.2.5. Section 4

You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

D.2.6. Section 5

You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

D.2.7. Section 6

Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these

terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

D.2.8. Section 7

If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

D.2.9. Section 8

If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

D.2.10. Section 9

The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

D.2.11. Section 10

If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

D.2.12. NO WARRANTY Section 11

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

D.2.13. Section 12

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

D.3. How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.> Copyright (C) <year> <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type “show w”. This is free software, and you are welcome to redistribute it under certain conditions; type “show c” for details.

The hypothetical commands “show w” and “show c” should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than “show w” and “show c”; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program “Gnomovision” (which makes passes at compilers) written by James Hacker.

<signature of Ty Coon>, 1 April 1989 Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.